

PTIME Computation of Transitive Closures of Octagonal Relations

Filip Konečný

Swiss Federal Institute of Technology Lausanne (EPFL)

Abstract. Computing transitive closures of integer relations is the key to finding precise invariants of integer programs. In this paper, we study *difference bounds* and *octagonal* relations and prove that their transitive closure is a PTIME-computable formula in the existential fragment of Presburger arithmetic. This result marks a significant complexity improvement, as the known algorithms have EXPTIME worst case complexity.

1 Introduction

This paper gives the first polynomial-time algorithm for computing *closed forms* of *difference bounds* and *octagonal* relations. Difference bounds (DB) relations are relations defined as conjunctions over atomic propositions of the form $x - y \leq c$ where c is an integer and x, y range over unprimed and primed variables $\mathbf{x} \cup \mathbf{x}'$. Octagonal relations generalize difference bounds relation by allowing conjuncts of the form $\pm x \pm y \leq c$. Both classes of relations are widely used as domains in verification of software and hardware [11, 12].

Given a binary relation R on states (represented as a formula with primed and unprimed variables) a *closed form* R is another formula $\hat{R}(k)$ containing primed and unprimed variables as well as a parameter variable k , such that substituting the parameter k with any integer $n \geq 1$ gives a precise description of R^n , the n -th power of R . The main result of this paper is a polynomial-time algorithm that, given the formula R in the form of octagonal constraints computes a closed form $\hat{R}(k)$ as a formula in the existential fragment of Presburger arithmetic. This result immediately extends to the computation of an expression for transitive closure, because $R^+ \Leftrightarrow \exists k \geq 1 . \hat{R}(k)$.

Approaches for computing the precise closed form of iterated relation compositions are referred to as *acceleration* algorithms. Known acceleration algorithms for the two classes of relations are based on the notion of *periodicity* and compute closed forms of the size that is polynomial in the size of the *prefix* and the *period* of a relation, both of which can be exponentially large in the binary size of a relation $\|R\|_2$. Intuitively, n -th power of a DB relation R can be obtained by computing minimal weights of paths between pairs of vertices in certain graphs (called *unfolded constraint graphs* of R and denoted G_R^n). For a fixed pair, minimal weights evolve periodically as a function of n . Due to these exponential bounds, an algorithm for computing closed forms that runs in time that is polynomial in $\|R\|_2$ must necessarily be based on a method different than explicitly computing periodicity. This paper presents the first such algorithm.

Overview First, we study difference bounds relations (Section 3 gives a background). Our main observation is that the problem of computing a closed form of a DB relation R can be reduced to the computation of closed forms of two PTIME-computable DB relations R_{fw} and R_{bw} such that $R_{fw} (R_{bw})$ belongs to a fragment called *forward (backward) one-directional* DB relations which contains DB relations of the form

$$\bigwedge_{ij} x_i - x'_j \leq c_{ij} \quad \left(\bigwedge_{ij} x'_i - x_j \leq c_{ij}, \text{ respectively} \right)$$

We first study these (dual) fragments and give a PTIME algorithm which computes the closed form in the existential fragment of Presburger arithmetic (Section 4). The main insight of this algorithm is that the closed form can be defined by encoding polynomially many *path schemes* which can be thought of as regular patterns that capture all paths with minimal weight in unfolded constraint graphs.

Next, we observe that for a fixed pair of vertices (u, v) in an unfolded constraint graph, any path ρ from u to v can be *normalized*, i.e. replaced with another path ρ' from u to v such that the weight of ρ' is not greater than the weight of ρ and ρ' is in a normal form (Section 5).

Then, we define the relations R_{fw} and R_{bw} and show that there exists an integer B of polynomial size such that every normalized path ρ in \mathcal{G}_R^{2B+n} can be viewed as a concatenation of several paths from $\mathcal{G}_{R_{fw}}^n$, $\mathcal{G}_{R_{bw}}^n$ and \mathcal{G}_R^B (Section 6). Since paths from $\mathcal{G}_{R_{fw}}^n$ and $\mathcal{G}_{R_{bw}}^n$ are captured by closed forms $\hat{R}_{fw}(n)$ and $\hat{R}_{bw}(n)$ (both PTIME-computable), and since paths in \mathcal{G}_R^B are captured by R^B (also PTIME-computable, since B is polynomially large), it follows that $\hat{R}_{fw}(n)$, $\hat{R}_{bw}(n)$, and R^B can be combined to form a closed form $\hat{R}(2B+n)$.

Finally, in Section 7, we show that these methods and results can be generalized to compute closed forms of octagonal relations in polynomial time as well. Section 8 concludes.

Related work Octagonal constraints [11] are well known in abstract interpretation as an abstract domain for over-approximating sets of reachable states. Transitive closure algorithms for octagonal relations [2] are the core of reachability analysis techniques based on computation of procedure summaries [9] or on accelerated interpolation [8].

DB and octagonal relations have been shown to have Presburger definable transitive closures [7,6,1] and to have periodic characterization [2]. An algorithm from [2] computes a transitive closure whose size is polynomial in the binary size of the relation $\|R\|_2$ and in the size of the prefix and period. Since relations whose prefix or period increases exponentially in $\|R\|_2$ can be constructed, the exponential lower bound on the size of the computed transitive closure follows.

Recently, [5] proves that both prefix and period can also be upper-bounded by a single exponential and moreover, shows NP-completeness of the reachability problem for *flat counter systems*, a class of integer programs without nested loops where each loop (non-loop) transition is described by an octagonal relation (QFPA¹ formula). Moreover,

¹ Quantifier-Free Presburger Arithmetic

[5] presents a non-deterministic reduction to satisfiability of QFPA formulas (an NP-complete problem), essentially by first guessing the prefix and period and then guessing one of exponentially many disjuncts of the transitive closure, for each loop. Our present result can turn this reduction into a deterministic one, since we can directly compute the transitive closure of each loop.

2 Preliminary Definitions

In the rest of this paper, let $N \geq 1$ and let $\mathbf{x} = \{x_1, x_2, \dots, x_N\}$ be a set of variables ranging over \mathbb{Z} . For each $n \in \mathbb{Z}$, we define a fresh copy of variables $\mathbf{x}^{(n)} \stackrel{\text{def}}{=} \{x_1^{(n)}, \dots, x_N^{(n)}\}$. Similarly, \mathbf{x}' denotes a fresh copy of primed variables $\mathbf{x}' = \{x'_1, \dots, x'_N\}$. We assume that the reader is familiar with Presburger arithmetic (PA). For a PA formula ϕ , let $\text{atoms}(\phi)$ denote the set of atomic propositions in ϕ , and $\phi[t/x]$ denote the formula obtained by substituting the variable x with the term t . $\text{card}(S)$ denotes the cardinality of a set S and $\text{abs}(c)$ denotes the absolute value of $c \in \mathbb{Z}$. A *valuation* of \mathbf{x} is a function $\mathbf{v} : \mathbf{x} \rightarrow \mathbb{Z}$. The set of all such valuations is denoted by $\mathbb{Z}^{\mathbf{x}}$. Given a relation $R \subseteq \mathbb{Z}^{\mathbf{x}} \times \mathbb{Z}^{\mathbf{x}}$, we denote by R^i , for $i > 0$, the i -times composition of R with itself. We denote by $R^+ = \bigcup_{i=1}^{\infty} R^i$ the *transitive closure* of R . If $R(\mathbf{x}, \mathbf{x}')$ defines R , we denote by $R^n(\mathbf{x}, \mathbf{x}')$ a formula that defines the n -th power R^n . A *closed form* of R is a formula $\hat{R}(k, \mathbf{x}, \mathbf{x}')$, where $k \notin \mathbf{x}$, such that $\hat{R}[n/k]$ defines R^n , for all $n \geq 1$. For a weighted graph G and a pair of vertices u, v , we denote by $\text{min-weight}(u, v, G)$ the minimal weight over all paths from u to v in G .

3 Difference Bounds Relations

Definition 1. A formula $\phi(\mathbf{x})$ is a difference bounds (DB) constraint if it is a finite conjunction of atomic propositions of the form $x_i - x_j \leq \alpha_{ij}$, $1 \leq i, j \leq N$, where $\alpha_{ij} \in \mathbb{Z}$. A relation $R \subseteq \mathbb{Z}^{\mathbf{x}} \times \mathbb{Z}^{\mathbf{x}}$ is a difference bounds relation if it can be defined by a difference bounds constraint $\phi_R(\mathbf{x}, \mathbf{x}')$.

Difference bounds constraints are represented as graphs. If $\phi(\mathbf{x})$ is a difference bounds constraint, then *constraint graph* $G_\phi = \langle \mathbf{x}, \rightarrow \rangle$ is a weighted graph, where each vertex corresponds to a variable, and there is an edge $x_i \xrightarrow{\alpha_{ij}} x_j$ in G_ϕ if and only if there exists a constraint $x_i - x_j \leq \alpha_{ij}$ in ϕ (Fig. 1(a)). The following result on existential quantification is well known [10]:




Fig. 1. The constraint graph G_R and its 3-times unfolding G_R^3 for a difference bounds relation $R \Leftrightarrow x_2 - x'_1 \leq -1 \wedge x_3 - x'_2 \leq 0 \wedge x_1 - x'_3 \leq 0 \wedge x'_4 - x_4 \leq 0 \wedge x'_4 - x_4 \leq 0$.

Proposition 1. Let $\phi(x_0, \mathbf{x})$ be a DB constraints. Then, $R(x_0, \mathbf{x})$ is consistent if and only if G_ϕ contains no negative cycle. If $\phi(x_0, \mathbf{x})$ is consistent, then

$$\exists x_0 . \phi(x_0, \mathbf{x}) \Leftrightarrow \bigwedge_{x, y \in \mathbf{x}} x - y \leq \text{min-weight}(x, y, G_\phi)$$

Moreover, consistency check and computation of $\exists x_0. \phi(x_0, \mathbf{x})$ is in $O(\|R\|_2)$ time.

Consequently, DB relations are closed under relational composition, i.e. $R^n(\mathbf{x}, \mathbf{x}')$ is a DB constraint for all $n \geq 1$. The n -th power of R can be seen as a constraint graph consisting of n copies of \mathcal{G}_R (see Fig. 1(b)):

Definition 2. Let $n \geq 1$ and $R(\mathbf{x}, \mathbf{x}')$ be a DB constraint. Then, the n -times unfolding of \mathcal{G}_R is defined as $\mathcal{G}_R^n \stackrel{\text{def}}{=} \bigcup_{i=0}^{n-1} \mathcal{G}_{R(\mathbf{x}^{(i)}, \mathbf{x}^{(i+1)})}$.

The vertices $\mathbf{x}^{(0)} \cup \mathbf{x}^{(n)}$ of \mathcal{G}_R^n are called *extremal*. A path in \mathcal{G}_R^n is said to be *extremal* if its first and last vertex are both extremal. The next lemma gives means to compute $R^n(\mathbf{x}, \mathbf{x}')$ and test its consistency, by analyzing extremal paths of \mathcal{G}_R^n .

Lemma 1. (Lemma 6 in [4]) Let $n \geq 1$ and $R(\mathbf{x}, \mathbf{x}')$ be a DB constraints. Then, $R^n(\mathbf{x}, \mathbf{x}')$ is consistent if and only if \mathcal{G}_R^n contains no extremal cycle with negative weight. If $R^n(\mathbf{x}, \mathbf{x}')$ is consistent, then $R^n(\mathbf{x}, \mathbf{x}')$ can be computed as

$$\bigwedge_{1 \leq i, j \leq N} x_i - x_j \leq \min\text{-weight}(x_i^{(0)}, x_j^{(0)}, \mathcal{G}_R^n) \wedge x'_i - x'_j \leq \min\text{-weight}(x_i^{(n)}, x_j^{(n)}, \mathcal{G}_R^n) \wedge \\ x_i - x'_j \leq \min\text{-weight}(x_i^{(0)}, x_j^{(n)}, \mathcal{G}_R^n) \wedge x'_i - x_j \leq \min\text{-weight}(x_i^{(n)}, x_j^{(0)}, \mathcal{G}_R^n)$$

Moreover, consistency check and computation of $R^n(\mathbf{x}, \mathbf{x}')$ is in $O(\|R\|_2 \cdot \log_2 n)$ time.

Paths in Unfoldings of \mathcal{G}_R In this paper, when the exact number of iterations does not matter, we sometimes consider paths in the bi-infinite unfolding ${}^\infty\mathcal{G}_R$ of \mathcal{G}_R , defined as

$${}^\infty\mathcal{G}_R \stackrel{\text{def}}{=} \bigcup_{i \in \mathbb{Z}} \mathcal{G}_{R(\mathbf{x}^{(i)}, \mathbf{x}^{(i+1)})}$$

Note that each edge in ${}^\infty\mathcal{G}_R$ is either *forward* (i.e. of the form $x_i^{(p)} \xrightarrow{\alpha} x_j^{(p+1)}$ for some $1 \leq i, j \leq N$ and $p, \alpha \in \mathbb{Z}$), *backward* ($x_i^{(p+1)} \xrightarrow{\alpha} x_j^{(p)}$), or *vertical* ($x_i^{(p)} \xrightarrow{\alpha} x_j^{(p)}$). A *path* is a sequence of the form (see Fig. 2 for illustrations)

$$\rho = x_{i_0}^{(p_0)} \xrightarrow{\alpha_0} x_{i_1}^{(p_1)} \xrightarrow{\alpha_1} \dots \xrightarrow{\alpha_{n-1}} x_{i_n}^{(p_n)}$$

for some $n \geq 0$ where $x_{i_k}^{(p_k)} \xrightarrow{\alpha_k} x_{i_{k+1}}^{(p_{k+1})}$ is an edge in ${}^\infty\mathcal{G}_R$, for each $0 \leq k < n$. We say that a variable x_{i_k} *occurs* on ρ at *position* p_k , for each $0 \leq k \leq n$. We say that ρ is *forward* (*backward*, *vertical*) if $p_0 < p_n$ ($p_0 > p_n$, $p_0 = p_n$, respectively). The *length* and *relative length* of ρ is defined as $|\rho| = n$ and $\|\rho\| = \text{abs}(p_n - p_0)$. The *weight* of ρ is defined as $\omega(\rho) = \alpha_0 + \dots + \alpha_{n-1}$. We write $\text{vars}(\rho)$ for the set $\{x_{i_0}, \dots, x_{i_n}\}$, $\text{positions}(\rho)$ for the set $\{p_0, \dots, p_n\}$, and $\text{vertices}(\rho)$ for the set $\{x_{i_0}^{(p_0)}, \dots, x_{i_n}^{(p_n)}\}$. We say that ρ is *repeating* if $p_0 \neq p_n$ and $i_0 = i_n$. We say that ρ is *elementary* if all vertices $x_{i_0}^{(p_0)}, \dots, x_{i_n}^{(p_n)}$ are distinct, with the exception of $x_{i_0}^{(p_0)}$ and $x_{i_n}^{(p_n)}$, which might be equal. We say that ρ is *essential* if all variables x_{i_0}, \dots, x_{i_n} are distinct, with the exception of x_{i_0} and x_{i_n} , which might be equal. Clearly, each essential path is also elementary. Note that the




Fig. 2. Paths in unfolded constraint graphs. Path π_1 is repeating and essential, π_2 is repeating, π_3 is essential, π_4 is neither essential nor elementary, π_5 is essential. We have $|\pi_1| = 4$, $|\pi_2| = 3$, $\|\pi_2\| = 3$, $\|\pi_1\| = \|\pi_3\| = \|\pi_4\| = 2$, $\|\pi_5\| = \|\pi_6\| = \|\pi_7\| = 0$, $\text{positions}(\pi_6) = \{3, 4, 5\}$, $\text{vars}(\pi_2) = \{x_4, x_5\}$, $\text{vertices}(\pi_2) = \{x_4^{(0)}, x_5^{(1)}, x_5^{(2)}, x_4^{(3)}\}$.

length of an essential path is bounded by N . A *subpath* of ρ is any path of the form $x_{i_a}^{(p_a)} \rightarrow \dots \rightarrow x_{i_b}^{(p_b)}$ where $0 \leq a \leq b \leq n$. We denote by $\vec{\rho}^{(k)} : x_{i_0}^{(p_0+k)} \rightarrow \dots \rightarrow x_{i_n}^{(p_n+k)}$ the path obtained by *shifting* ρ by k , where $k \in \mathbb{Z}$. A path ρ is said to be *isomorphic* with another path ρ' if and only if $\rho' = \vec{\rho}^{(k)}$, for some $k \in \mathbb{Z}$. Consider a path $\pi = x_{j_0}^{(q_0)} \xrightarrow{\beta_0} \dots \xrightarrow{\beta_{m-1}} x_{j_m}^{(q_m)}$. The *concatenation* $\rho.\pi$ is defined if $x_{i_n}^{(p_n)} = x_{j_0}^{(q_0)}$. If $i_n = j_0$, we write $\rho.\pi$ as a shorthand for $\rho.\vec{\pi}^{(p_n-q_0)}$. If ρ is repeating and $k \geq 1$, we define the *k-th power* of ρ as the k -times concatenation of ρ with itself, e.g. $\rho^3 = \rho.\rho.\rho$. We next define the notion of a *compatible* path.

Definition 3. Let ρ, ρ' be paths in \mathcal{G}_R^n for some $n \geq 1$. We say that ρ' is compatible with ρ (denoted $\rho' \preceq \rho$) if and only if (i) both ρ and ρ' are of the form $x_i^{(k)} \rightarrow \dots \rightarrow x_j^{(\ell)}$ for some $1 \leq i, j \leq n$ and $0 \leq k, \ell \leq n$, and (ii) $\omega(\rho') \leq \omega(\rho)$.

Balanced relations We say that a difference bounds constraint $R(\mathbf{x}, \mathbf{x}')$ is *balanced* whenever $(x - y \leq c) \in \text{atoms}(R)$ if and only if $(x' - y' \leq c) \in \text{atoms}(R)$. Note that the relation R_b (called the *balanced closure* of R) defined below is balanced:

$$R_b \stackrel{\text{def}}{=} R \wedge \bigwedge_{(x-y \leq c) \in \text{atoms}(R)} x' - y' \leq c \wedge \bigwedge_{(x'-y' \leq c) \in \text{atoms}(R)} x - y \leq c$$

We next show that the computation of the closed form for R can be reduced to the computation of the closed form of its balanced closure:

Proposition 2. Let $R(\mathbf{x}, \mathbf{x}')$ be a DB constraint, $R_b(\mathbf{x}, \mathbf{x}')$ be its balanced closure, and $\widehat{R}_b(\ell, \mathbf{x}, \mathbf{x}')$ be the closed form of R_b . Then, $\widehat{R}(k, \mathbf{x}, \mathbf{x}')$ can be defined as:

$$\bigvee_{i=1}^2 (k = i \wedge R^i(\mathbf{x}, \mathbf{x}')) \vee \exists \mathbf{x}_1, \mathbf{x}_2. k \geq 3 \wedge R(\mathbf{x}, \mathbf{x}_1) \wedge \widehat{R}_b(\ell, \mathbf{x}_1, \mathbf{x}_2)[k - 2/\ell] \wedge R(\mathbf{x}_2, \mathbf{x}')$$

In the following sections, we study balanced DB relations. Finally, as a consequence of Proposition 2, we show that our results can be generalized to arbitrary DB relations.

4 Closed Forms for One-directional Difference Bounds Relations

We say that a DB constraint $R(\mathbf{x}, \mathbf{x}')$ is *one-directional* if it is either (i) a conjunction of the form $\bigwedge_{ij} x_i - x'_j \leq c_{ij}$ (*forward* one-directional) or (ii) a conjunction of the form $\bigwedge_{ij} x'_i - x_j \leq c_{ij}$ (*backward* one-directional). Clearly, the two cases are dual: R is forward one-directional if and only if its inverse R^{-1} (which can be defined as $R(\mathbf{x}, \mathbf{x}')[\mathbf{x}'/\mathbf{x}, \mathbf{x}/\mathbf{x}']$) is backward one-directional. Consequently, a closed form of R can be directly obtained from a closed form of R^{-1} as:

$$\widehat{R}(k, \mathbf{x}, \mathbf{x}') \Leftrightarrow \widehat{R^{-1}}(k, \mathbf{x}, \mathbf{x}')[\mathbf{x}/\mathbf{x}', \mathbf{x}'/\mathbf{x}]$$

We can thus consider, without loss of generality, only forward one-directional relations. Let R be such relation. Clearly, \mathcal{G}_R^n contains only forwards edges for all $n \geq 1$. Hence, $|\rho| = \|\rho\|$ for each path ρ in \mathcal{G}_R^n and moreover, \mathcal{G}_R^n contains no cycle and R^n is thus consistent, for all $n \geq 1$. Then, by Proposition 1, computation of $R^n(\mathbf{x}, \mathbf{x}')$ amounts to computing, for each $1 \leq i, j \leq N$, the minimal weight over all paths in \mathcal{G}_R^n of the form $x_i^{(0)} \rightarrow \dots \rightarrow x_j^{(n)}$. We next show that minimal weight paths have, without loss of generality, regular shape in the sense that they are instances of *biquadratic path schemes*:

Definition 4. If σ, σ' are paths and λ is an empty or an essential repeating path such that $\sigma.\lambda.\sigma'$ is a non-empty path, the expression $\theta = \sigma.\lambda^*.\sigma'$ is called a *path scheme*. A path scheme encodes the infinite set of paths $[\![\theta]\!] = \{\sigma.\lambda^n.\sigma' \mid n \geq 0\}$. We say that θ is *biquadratic* if $|\sigma.\sigma'| \leq N^4$.

The following result is a consequence of Lemma 3 in [4]:

Lemma 2. Let R be a one-directional DB relation, let $n \geq 1$, and let ρ be an extremal path in \mathcal{G}_R^n . Then, there exists a compatible path ρ' and a biquadratic path scheme $\sigma.\lambda^*.\sigma'$, such that $\rho' \in [\![\sigma.\lambda^*.\sigma']\!]$.

By Lemma 2, minimal weight paths can be captured by a set Π of all biquadratic path schemes. For each such scheme $\sigma.\lambda^*.\sigma' \in \Pi$, we have $|\sigma.\sigma'| \leq N^4$ (by Def. 4) and $|\lambda| \leq N$ (since the length of essential paths is bounded by N). In the worst case, each vertex of \mathcal{G}_R^n has N successors and hence, there are up to N^n paths in \mathcal{G}_R^n of the form $x_i^{(0)} \rightarrow \dots \rightarrow x_j^{(n)}$, for a fixed $1 \leq i, j \leq N$. Consequently, $\text{card}(\Pi)$ is of the order $2^{O(N)}$. We next show that it is sufficient to consider only polynomially many representants from Π . We first partition Π into polynomially many equivalence classes. Each class is determined by (i) first and last variables of σ , λ , and σ' , and by (ii) the length of λ and $\sigma.\sigma'$. Formally, the partition is defined as:

$$\Xi \stackrel{\text{def}}{=} \{\Pi_{ijkpq} \mid 1 \leq i, j, k \leq N, 0 \leq p \leq N^4, 0 \leq q \leq N, p+q > 0\}$$

where each $\Pi_{ijkpq} \subseteq \Pi$ is defined as follows: $\sigma.\lambda^*.\sigma' \in \Pi_{ijkpq}$ if and only if σ, λ, σ' are paths of the form:

$$\lambda = x_k^{(0)} \rightarrow \dots \rightarrow x_k^{(p)} \tag{1}$$

$$\sigma = x_i^{(0)} \rightarrow \dots \rightarrow x_k^{(r)}, \quad \sigma' = x_k^{(r)} \rightarrow \dots \rightarrow x_j^{(q)}, \quad \text{for some } 0 \leq r \leq q \tag{2}$$

Intuitively, p (q) determines the length of λ ($\sigma.\sigma'$) and k determines the variable on which λ connects with σ and σ' . Clearly, $\text{card}(\Xi)$ is of the order $O(N^8)$.

Let us fix i, j, k, p, q and assume that $\Pi_{ijkpq} \neq \emptyset$. It is easy to see that if there exists a path λ of the form (1), then there exists one with minimal weight. Similarly, if there exists a path $\sigma.\sigma'$ of the form (2), then there exists one with minimal weight. We define θ_{ijkpq} as the path scheme $\sigma.\lambda^*.\sigma'$ where λ and $\sigma.\sigma'$ are the minimal paths. It is easy to see that θ_{ijkpq} is *minimal* in Π_{ijkpq} in the following sense: $\omega(\sigma.\lambda^n.\sigma') \leq \omega(v.\mu^n.v')$ for each $v.\mu^n.v' \in \Pi_{ijkpq}$ and each $n \geq 0$. Hence, we can use θ_{ijkpq} as a representant of Π_{ijkpq} . The minimal representants can be computed in polynomial time:

Lemma 3. *The set $\{\theta_{ijkpq} \mid \Pi_{ijkpq} \neq \emptyset\}$ can be computed in PTIME.*

Next, we fix $1 \leq i, j \leq N$ and define:

$$S_{ij} \stackrel{\text{def}}{=} \{(|\sigma.\sigma'|, |\lambda|, \omega(\sigma.\sigma'), \omega(\lambda)) \mid \exists k, p, q. \theta_{ijkpq} = \sigma.\lambda^*.\sigma'\}$$

It follows from the previous arguments that S_{ij} represents all bi-quadratic schemes which capture paths from $x_i^{(0)}$ to $x_j^{(n)}$ and moreover, S_{ij} can be computed in polynomial time and its cardinality is polynomial. Then, the closed form of the sequence $\{\text{min-weight}(x_i^{(0)}, x_j^{(n)}, \mathcal{G}_R^n)\}_{n \geq 1}$ can be defined as:

$$\phi_{ij}(n, x_i, x'_j) \Leftrightarrow n \geq 1 \wedge \bigwedge_{(p,q,a,b) \in S_{ij}} \forall \ell. (\ell \geq 0 \wedge n = p + q \cdot \ell \Rightarrow (x_i - x'_j \leq a + b \cdot \ell))$$

Intuitively, each conjunct encodes a constraint of one scheme $\sigma.\lambda^*.\sigma'$: whenever the scheme captures a path of length n (i.e. $n = |\sigma.\sigma'| + \ell \cdot |\lambda| = p + q \cdot \ell$ where $\ell \geq 0$), the difference $x_i - x'_j$ must be upper-bounded by the corresponding weight $\omega(\sigma.\sigma') + \ell \cdot \omega(\lambda) = a + \ell \cdot b$. Equivalently, we can write:

$$\phi_{ij}(n, x_i, x'_j) \Leftrightarrow n \geq 1 \wedge \bigwedge_{(p,q,a,b) \in S_{ij}} (n \geq p \wedge q \mid n - p) \Rightarrow q \cdot (x_i - x'_j) \leq q \cdot a + b \cdot (n - p) \quad (3)$$

Then, we can define the closed form of R as:

$$\widehat{R}(n, \mathbf{x}, \mathbf{x}') \Leftrightarrow \bigwedge_{1 \leq i, j \leq N} \phi_{ij}(n, x_i, x'_j) \quad (4)$$

Clearly, $\phi_{ij}(n, x_i, x'_j)$ (and hence also $\widehat{R}(n, \mathbf{x}, \mathbf{x}')$) is a formula in the existential fragment of PA and of polynomial size, since $\text{card}(S_{ij})$ is polynomial. Thus, it follows from Lemma 3 that the whole computation of $\widehat{R}(n, \mathbf{x}, \mathbf{x}')$ is polynomial.

Theorem 1. *Let $R(\mathbf{x}, \mathbf{x}')$ be a one-directional DB constraint. Then, its closed form $\widehat{R}(n, \mathbf{x}, \mathbf{x}')$ can be computed in PTIME as a formula in the existential fragment of PA.*

5 Normalization of Paths in the Unfolded Constraint Graph

In this section, we consider only balanced DB relations and show that every extremal path in an unfolded constraint graph can be *normalized*. Intuitively, a path p from \mathcal{G}_R^n is

normalized if none of its subpaths that traverses only positions in the range $\{N^2, \dots, n - N^2\}$ is a *long corner*. Informally, a *corner* is a vertical path that stays either on the right or on the left side of the initial position. A corner is *long* if the distance between its minimal and maximal position exceeds the bound N^2 .

Definition 5. (Corners) Let ρ be a vertical path of the form $\rho = x_{i_0}^{(k_0)} \rightarrow \dots \rightarrow x_{i_m}^{(k_m)}$ for some $m \geq 1$ such that $k_0 = k_m$. If $\text{positions}(\rho) = \{k_0, \dots, k_0 + d\}$ for some $d \geq 0$, we say that ρ is a right corner of extent d . If $\text{positions}(\rho) = \{k_0 - d, \dots, k_0\}$ for some $d \geq 0$, we say that ρ is a left corner of extent d . We say that ρ is a corner if it is either a left corner or a right corner. We denote the extent of a corner ρ by $\text{extent}(\rho)$. We say that a corner ρ is basic if $k_0 \notin \{k_1, \dots, k_{m-1}\}$. We say that a corner ρ is long if $\text{extent}(\rho) > N^2$. We say that ρ is a lb-corner if it is both long and basic.

E.g., consider vertical paths from Fig. 1(c), where π_6 is a right corner, π_7 is a right basic corner, and π_5 is not a corner. Both π_6 and π_7 are short, since $\text{extent}(\pi_6) = \text{extent}(\pi_7) = 2 \leq N^2 = 5^2$. In the following, $\text{lb-corners}(\rho)$ ($\text{l-corners}(\rho)$, respectively) denotes the set of subpaths of ρ which are lb-corners (long corners, respectively). It is easy to show that if a path contains no lb-corner, it also contains no long corner. We are now ready to formalize the notion of a *normalized* path.

Definition 6. (Normalized paths) Let $n \geq 1$ and let ρ be an extremal path in \mathcal{G}_R^n . We say that ρ is normalized if none of its subpaths θ such that $\text{positions}(\theta) \subseteq \{N^2, \dots, n - N^2\}$ is a long corner.

E.g., the path in Fig. 3(a) is not normalized, due to the long corner θ .

Normalization We next give a high level idea of normalization. Given an integer $n \geq 1$ and an extremal path ρ_1 from \mathcal{G}_R^n , we construct a finite sequence $\{\rho_k\}_{k=1}^m$ of paths from \mathcal{G}_R^n for some $m \geq 1$ such that ρ_m is normalized and ρ_{k+1} is compatible with ρ_k , for each $1 \leq k \leq m-1$. By transitivity, we have that ρ_m is compatible with ρ_1 . For each $1 \leq k < m$, the path ρ_{k+1} is obtained from ρ_k by *substituting* some of its subpaths with a compatible path.

Definition 7. (Substitution) If $\rho.\theta.\rho'$ and θ' are paths in \mathcal{G}_R^n such that $\theta' \preceq \theta$, the substitution of θ in $\rho.\theta.\rho'$ with θ' is defined as $(\rho.\theta.\rho')[\theta'/\theta] \stackrel{\text{def}}{=} \rho.\theta'.\rho'$.

The subpaths of $\rho_1, \dots, \rho_{m-1}$ that are substituted are certain paths called *segments*: Informally, *segments* of ρ are the unique subpaths of ρ that traverse only positions $\{p, \dots, q\}$ for some fixed parameters $p \leq q$.

Definition 8. (Path segments) Let $n \geq 1$ and let ρ be a path in \mathcal{G}_R^n . Let $0 \leq p \leq q \leq n$ be integers and let \mathcal{H} be the (unique) subgraph of \mathcal{G}_R^n obtained by removing every edge τ such that $\text{positions}(\tau) \subseteq \{p, \dots, q\}$. We define $\text{segments}(\rho, p, q)$ to be the (unique) sequence ξ_1, \dots, ξ_m , for some $m \geq 0$, such that

- ξ_i is a subpath of ρ such that $\text{positions}(\xi_i) \subseteq \{p, \dots, q\}$, for each $1 \leq i \leq m$
- there exist (unique) paths $\sigma_1, \dots, \sigma_{m+1}$ in \mathcal{H} such that $\rho = \sigma_1.\xi_1 \dots \sigma_m.\xi_m.\sigma_{m+1}$

In the rest of this paper, we write $\text{segments}(\rho)$ as a shorthand for $\text{segments}(\rho, N^2, n - N^2)$. If $\xi \in \text{segments}(\rho)$, we say that ξ is a *segment* of ρ . It is easy to verify that each segment of an extremal path ρ in G_R^n is of the form $\xi = x_i^{(p)} \rightarrow \dots \rightarrow x_j^{(q)}$ for some $1 \leq i, j \leq N$ and $p, q \in \{N^2, n - N^2\}$. As an example, consider the path in Fig. 3(a), which has one segment $\gamma.\theta.\gamma'$. The next proposition allows us to use an alternative characterization of normalized paths:

Proposition 3. *Let $n \geq 1$ and let ρ be an extremal path in G_R^n . Then, ρ is normalized if and only if $\text{lb-corners}(\xi) = \emptyset$ for each $\xi \in \text{segments}(\rho)$.*

Termination argument We argue that the sequence $\{\rho_k\}_{k=1}^m$ is finite, by tracking, for each segment, the distance of the *first lb-corner* from the end of the segment:

Proposition 4. (Finding the first lb-corner) *Let ρ be a path such that $\text{lb-corners}(\rho) \neq \emptyset$. Then, ρ has subpaths ρ_1, θ, ρ_2 such that $\rho = \rho_1.\theta.\rho_2$, $\text{lb-corners}(\rho_1.\theta) = \{\theta\}$, and $\text{extent}(\theta) = N^2 + 1$. The corner θ is called the *first lb-corner* of ρ .*

E.g., θ is the first lb-corner of the path $\sigma_1.\gamma.\theta.\gamma'.\sigma_2$ in Figure 3(a). We define $\text{lb-segments}(\rho)$ to be the subsequence of $\text{segments}(\rho)$ obtained by erasing every segment ξ such that $\text{lb-corners}(\xi) = \emptyset$. For each $1 \leq k < m$, we guarantee that if $\text{lb-segments}(\rho_k) = \xi_1, \dots, \xi_a$ for some $a \geq 1$, then either

$$\begin{aligned} (1) \text{ lb-segments}(\rho_{k+1}) &= \xi_2, \dots, \xi_a, \text{ or} \\ (2) \text{ lb-segments}(\rho_{k+1}) &= \zeta, \xi_2, \dots, \xi_a \text{ for some } \zeta \end{aligned} \tag{5}$$

and moreover, in case (2), ξ_1 and ζ are paths such that

- $\xi_1 = \gamma_1.\theta_1.\gamma'_1$ for some paths $\gamma_1, \theta_1, \gamma'_1$ and θ_1 is the first lb-corner of ξ_1 ,
- $\zeta = \gamma_2.\theta_2.\gamma'_2$ for some paths $\gamma_2, \theta_2, \gamma'_2$ and θ_2 is the first lb-corner of ζ , and
- $|\gamma'_2| < |\gamma'_1|$

Intuitively, ρ_k and ρ_{k+1} have the same segments with long corners, with the exception of one segment ξ_1 , which is either eliminated (case 1), or replaced with another segment ζ (case 2) such that the length of the unique suffix γ'_2 of ζ after its first corner is strictly smaller than the length of the unique suffix γ'_1 of ξ_1 after its first corner. Hence, the number of consecutive applications of the case 2 is bounded by $|\xi_1|$. Clearly, if $|\xi_1| = 0$, then only case 1 may happen, which decreases the number of segments with long corners, and therefore guarantees termination.

Transforming segments with long corners Let $n \geq 1$, ρ be an extremal path in G_R^n , and let $\xi \in \text{lb-segments}(\rho)$. We show how to construct a path ρ' that is compatible with ρ and moreover satisfies the termination properties from (5). By Proposition 4, there exists a unique corner θ such that $\xi = \gamma.\theta.\gamma'$, $\text{lb-corners}(\gamma.\theta) = \{\theta\}$, and $\text{extent}(\theta) = N^2 + 1$, for some paths γ, γ' . Fig. 3(a) depicts such situation. Suppose that ξ starts at position N^2 and ends at position $n - N^2$ (the other three cases are symmetric). Then, it is not difficult to show that θ is a right corner. The following lemma states a key result which allows us to either shorten or decompose the corner θ .

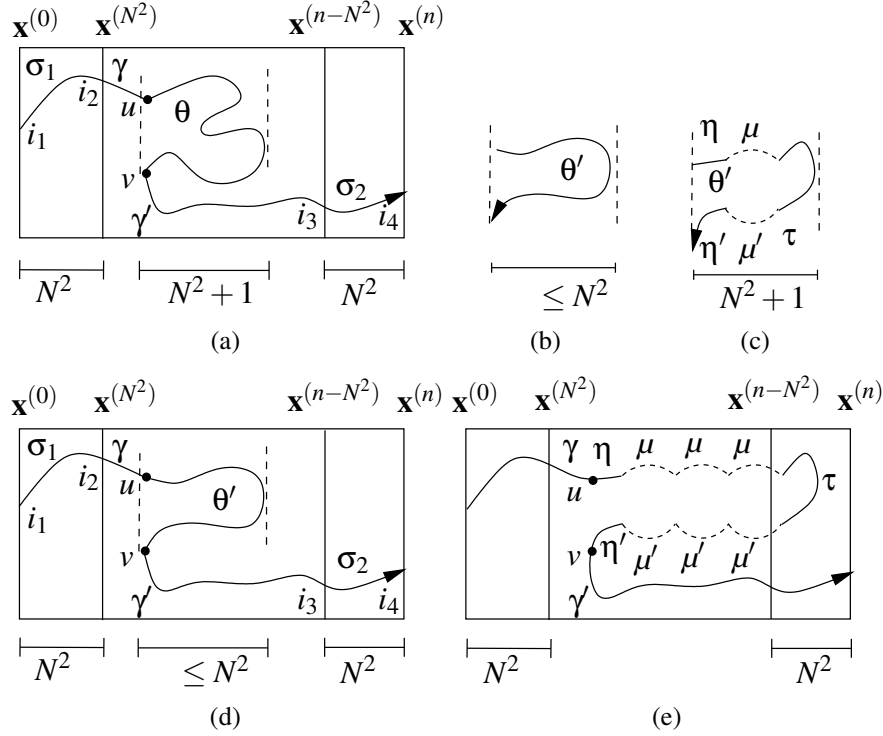


Fig. 3. Transformation of a segment with long corners

Lemma 4. (Corner shortening / decomposition) Let ρ be a right (left) lb-corner such that $\text{extent}(\rho) = N^2 + 1$. Then, there exists a compatible right (left) lb-corner ρ' such that either (i) $\text{extent}(\rho') \leq N^2$ or (ii) $\text{extent}(\rho') = N^2 + 1$ and ρ' has subpaths $\eta, \mu, \tau, \mu', \eta'$ such that $\rho' = \eta \cdot \mu \cdot \tau \cdot \mu' \cdot \eta'$, μ is a forward (backward) repeating path, μ' is a backward (forward) repeating path, τ is a right (left) corner, $\|\eta\| = \|\eta'\|$, $1 \leq \|\mu\| = \|\mu'\| \leq N^2$, and $\omega(\mu) + \omega(\mu') < 0$. Moreover, for all $k \geq 0$, $\eta \cdot \mu^k \cdot \tau \cdot \mu'^k \cdot \eta'$ is a right (left) corner and $\text{lb-corners}(\eta \cdot \mu^k) = \text{lb-corners}(\mu'^k \cdot \eta') = \emptyset$.

Let θ' be the corner obtained by applying Lemma 4. Fig. 3(b) depicts the case (i) and Fig. 3(c) the case (ii).

First, suppose that the case (i) of Lemma 4 applies. Then, one can define

$$\zeta \stackrel{\text{def}}{=} \xi[\theta'/\theta] = (\gamma \cdot \theta \cdot \gamma')[\theta'/\theta] = \gamma \cdot \theta' \cdot \gamma'$$

(see Fig. 3(d)) and prove that $\text{lb-corners}(\gamma \cdot \theta') = \emptyset$, by using the fact from Lemma 4 that $\text{extent}(\theta') \leq N^2$. If $\text{lb-corners}(\zeta) = \emptyset$, then the case 1 in (5) applies. If $\text{lb-corners}(\zeta) \neq \emptyset$, one can infer from $\text{lb-corners}(\gamma \cdot \theta') = \emptyset$ that the first corner of ζ involves at least one edge of γ' and hence, that the distance of the first lb-corner in ζ from the end of ζ strictly decreases, i.e. that the case 2 in (5) applies. Hence, the termination property

is preserved. We have $\theta' \preceq \theta$, by Lemma 4. Consequently, $\zeta \preceq \xi$ and we can define $\rho' \stackrel{\text{def}}{=} \rho[\zeta/\xi]$ and see that also $\rho' \preceq \rho$.

Second, suppose that the case (ii) of Lemma 4 applies. Let $\theta' = \eta.\mu.\tau.\mu'.\eta'$ be the decomposition of θ' given by Lemma 4. Note that μ, μ' are repeating and have the same relative length and opposite directions. Hence, we can define the following path ($\ell \geq 1$ is a parameter):

$$\bar{\theta} \stackrel{\text{def}}{=} \eta.\mu^\ell.\tau.\mu'^\ell.\eta'$$

We define $\rho' \stackrel{\text{def}}{=} \rho[\bar{\theta}/\theta']$. By Lemma 4, $\omega(\mu) + \omega(\mu') < 0$. Consequently, $\omega(\bar{\theta}) \leq \omega(\theta')$ for any $\ell \geq 1$ and hence, ρ' is compatible with ρ , i.e. $\rho' \preceq \rho$. By Lemma 4, $1 \leq \|\mu\| = \|\mu'\| \leq N^2$ and hence, one can choose ℓ sufficiently high and make the path $\bar{\theta}$ reach a position in the range $\{n - N^2 + 1, \dots, n\}$, formally: $n - N^2 + 1 \in \text{positions}(\bar{\theta})$. See Fig. 3(e) for an illustration. Thus, the segment ξ in ρ is replaced by two segments ζ', ζ in ρ' . Intuitively, ζ' has the subpath $\gamma.\eta.\mu^{\ell-1}$ and ζ has the subpath $\mu'^{\ell-1}.\eta'.\gamma'$. Next, we can apply the following property, which is by Lemma 4:

$$\text{lb-corners}(\eta.\mu^k) = \text{lb-corners}(\mu'^k.\eta') = \emptyset$$

to prove that $\text{lb-corners}(\zeta') = \emptyset$ and $\text{lb-corners}(\mu'^{\ell-1}.\eta') = \emptyset$. The former implies that ζ' is a segment with no long corners. The latter can then be used to prove that ζ has the same properties as ζ in the previous paragraph (for case (i)), i.e. that the termination properties are satisfied in this case as well.

We can thus conclude that every extremal path can be normalized.

Theorem 2. *Let $R(\mathbf{x}, \mathbf{x}')$ be a balanced DB constraint, let $n \geq 1$ be an integer, and let ρ be a path between extremal vertices of \mathcal{G}_R^n . Then, there exists a normalized path ρ' such that $\rho' \preceq \rho$.*

6 Closed Forms for Difference Bounds Relations

By Lemma 1, relation R^n is consistent if \mathcal{G}_R^n contains no extremal cycle with negative weight and moreover, consistent relation R^n can be defined as a conjunction of constraints each of which corresponds to a minimal extremal path. Hence, proving that a formula $\phi(\mathbf{x}, \mathbf{x}')$ defines R^n amounts to showing that $\phi(\mathbf{x}, \mathbf{x}')$ implies only those constraints represented by extremal paths in \mathcal{G}_R^n . Consequently, a closed form $\hat{R}(k, \mathbf{x}, \mathbf{x}')$ must satisfy the above for each $k \geq 1$. In this section, we show how to define such formula, in several steps. First, we strengthen the relation R in a way that enables us to shortcut every short corner with a single vertical edge (Section 6.1). Second, we define a formula that encodes paths that do not contain long corners (Section 6.2). Third, we generalize this encoding to extremal paths (Section 6.3), by exploiting the fact that such paths can be decomposed into segments according to Def. 8 and that segments of extremal normalized paths contain no long corners. Finally, we show how the formula that encodes extremal paths can be used to define a closed form (Section 6.4).

6.1 Shortcutting Short Corners

Consider the strengthened relation R_s in Fig. 4(a). We prove that for each $n \geq 1$, each short corner in \mathcal{G}_R^n has a compatible vertical edge in $\mathcal{G}_{R_s}^n$ (see Fig. 4(b-c)).




Fig. 4. Shortcutting a short corner by strengthening a relation

Proposition 5. (Eliminating Short Corners) Let $R(\mathbf{x}, \mathbf{x}')$ be a balanced DB constraint, let $n \geq 1$ be an integer and let θ be a short corner in \mathcal{G}_R^n of the form $x_i^{(k)} \rightarrow \dots \rightarrow x_j^{(k)}$. Then:

$$\begin{aligned} R_s(\mathbf{x}, \mathbf{x}') &\Rightarrow S_{fw}(\mathbf{x}) \Rightarrow x_i - x_j \leq \omega(\theta) && (\text{if } \theta \text{ is right}) \\ R_s(\mathbf{x}, \mathbf{x}') &\Rightarrow S'_{bw}(\mathbf{x}') \Rightarrow x'_i - x'_j \leq \omega(\theta) && (\text{if } \theta \text{ is left}) \end{aligned}$$

Consequently, there is a compatible vertical edge $x_i^{(k)} \xrightarrow{c} x_j^{(k)}$ in $\mathcal{G}_{R_s}^n$, for some $c \leq \omega(\theta)$.

The intuition is that if we view the above short right corner θ as an extremal path in $\mathcal{G}_R^{N^2}$ that starts at position 0, then we have, by Lemma 1, that $R^{N^2}(\mathbf{x}, \mathbf{x}') \Rightarrow x_i - x_j \leq \omega(\theta)$, and hence the first implication in Proposition 5 holds, by the definition of S_{fw} and R_s .

6.2 Encoding Paths without Long Corners

The strengthening from Section 6.1 can be used to *straighten* paths which have only short corners. Informally, a straightened path is either (i) a sequence of forward edges, or (ii) a sequence of backward edges, or (iii) a single vertical edge. Let ξ be an extremal path in \mathcal{G}_R^n such that $\text{l-corners}(\xi) = \emptyset$. First, suppose that ξ is forward. Then, ξ can be viewed as a sequence of forward edges and right corners in \mathcal{G}_R^n (Fig. 5(b)). By Proposition 5, each corner can be shortcut by a vertical edge, and hence we obtain an equivalent path ξ' in $\mathcal{G}_{R_s}^n$ which is a sequence of forward and vertical edges in $\mathcal{G}_{R_s}^n$ (Fig.




Fig. 5. Path straightening

5(c)). Then, every subpath of ξ' of the form $(\text{vertical-edge})^*.fw\text{-edge}$ can be replaced by a (transitively) implied forward edge in $\mathcal{G}_{R_{fw}}^n$ where R_{fw} is defined in Fig. 5(a). and thus obtaining an equivalent path ξ'' in $\mathcal{G}_{R_{fw}}^n$ that contains only forward edges (Fig. 5(d)). Then, ξ'' is encoded by $\widehat{R}_{fw}(\ell, \mathbf{x}, \mathbf{x}')[n/\ell]$ and hence also in $\phi(\ell, \mathbf{x}, \mathbf{x}')[n/\ell]$ defined as:

$$\phi(\ell, \mathbf{x}, \mathbf{x}') \Leftrightarrow \widehat{R}_{fw}(\ell, \mathbf{x}, \mathbf{x}') \wedge \widehat{R}_{bw}(\ell, \mathbf{x}, \mathbf{x}') \wedge S_{fw}(\mathbf{x}) \wedge S_{bw}(\mathbf{x}') \quad (6)$$

If ξ is an extremal right corner, then it is encoded by $S_{fw}(\mathbf{x})$ (by Proposition 5) and hence also by $\phi(\ell, \mathbf{x}, \mathbf{x}')$ (since S_{fw} is its conjunct). The other cases (backward extremal path, extremal left corner) are symmetric. Hence, $\phi(\ell, \mathbf{x}, \mathbf{x}')$ encodes all extremal paths in \mathcal{G}_R^n that have no long corners, in the following sense:

Proposition 6. (Encoding of paths without long corners) *Let $R(\mathbf{x}, \mathbf{x}')$ be a balanced DB constraint, let $n \geq 1$, and let ξ be an extremal path in \mathcal{G}_R^n , i.e. of the form $x_i^{(p)} \rightarrow^+ x_j^{(q)}$ for some $1 \leq i, j \leq N$ and $p, q \in \{0, n\}$. If $\text{l-corners}(\xi) = \emptyset$, then:*

1. $\phi(\ell, \mathbf{x}, \mathbf{x}')[n/\ell] \Rightarrow x_i - x'_j \leq \omega(\xi)$ if $p = 0, q = n$
2. $\phi(\ell, \mathbf{x}, \mathbf{x}')[n/\ell] \Rightarrow x'_i - x_j \leq \omega(\xi)$ if $p = n, q = 0$
3. $\phi(\ell, \mathbf{x}, \mathbf{x}')[n/\ell] \Rightarrow x_i - x_j \leq \omega(\xi)$ if $p = q = 0$
4. $\phi(\ell, \mathbf{x}, \mathbf{x}')[n/\ell] \Rightarrow x'_i - x'_j \leq \omega(\xi)$ if $p = q = n$

6.3 Encoding Extremal Paths

Consider the following formula (let \mathbf{y} and \mathbf{z} be fresh copies of variables in \mathbf{x}):

$$\psi(\ell, \mathbf{x}, \mathbf{x}') \Leftrightarrow \exists \mathbf{y}, \mathbf{z} . R^{N^2}(\mathbf{x}, \mathbf{y}) \wedge \phi(\ell, \mathbf{y}, \mathbf{z}) \wedge R^{N^2}(\mathbf{z}, \mathbf{x}') \quad (7)$$

We prove that for each $n \geq 1$, the formula $\psi(\ell, \mathbf{x}, \mathbf{x}')$ encodes every extremal path in \mathcal{G}_R^n , in the following sense.

Proposition 7. (Encoding of extremal paths) *Let $R(\mathbf{x}, \mathbf{x}')$ be a balanced DB constraint, let $n \geq 1$, and let ρ be an extremal normalized path in $\mathcal{G}_R^{2N^2+n}$. Then, ρ is of the form $x_i^{(p)} \rightarrow^+ x_j^{(q)}$ for some $1 \leq i, j \leq N$ and $p, q \in \{0, 2N^2 + n\}$ and moreover:*

1. $\psi(\ell, \mathbf{x}, \mathbf{x}')[n/\ell] \Rightarrow x_i - x'_j \leq \omega(\rho)$ if $p = 0, q = 2N^2 + n$
2. $\psi(\ell, \mathbf{x}, \mathbf{x}')[n/\ell] \Rightarrow x'_i - x_j \leq \omega(\rho)$ if $p = 2N^2 + n, q = 0$
3. $\psi(\ell, \mathbf{x}, \mathbf{x}')[n/\ell] \Rightarrow x_i - x_j \leq \omega(\rho)$ if $p = q = 0$
4. $\psi(\ell, \mathbf{x}, \mathbf{x}')[n/\ell] \Rightarrow x'_i - x'_j \leq \omega(\rho)$ if $p = q = 2N^2 + n$

The intuition behind the encoding is as follows. Let ρ be an extremal normalized path and let $\rho = \sigma_1.\xi_1 \dots \sigma_m.\xi_m.\sigma_{m+1}$ be its decomposition according to Def. 8. By Prop. 3, $\text{l-corners}(\xi_i) = \emptyset$ for each $1 \leq i \leq m$, and hence, by Prop. 6, ξ_i is encoded by $\phi(\ell, \mathbf{y}, \mathbf{z})$. For each $1 \leq i \leq m+1$, we have that σ_i is encoded in $R^{N^2}(\mathbf{x}, \mathbf{y})$ or in $R^{N^2}(\mathbf{z}, \mathbf{x}')$. Then, since $\rho = \sigma_1.\xi_1 \dots \sigma_m.\xi_m.\sigma_{m+1}$, one can show, by transitivity, that (7) encodes ρ . E.g.,

consider the path $\rho = \sigma_1.\gamma.\theta'.\gamma'.\sigma_2$ in Figure 3(d) and denote $\xi_1 = \gamma.\theta'.\gamma'$. Supposing ρ is normalized, we have:

$$\left(\begin{array}{l} R^{N^2}(\mathbf{x}, \mathbf{y}) \Rightarrow x_{i_1} - y_{i_2} \leq \omega(\sigma_1) \\ \phi(\ell, \mathbf{y}, \mathbf{z})[n/\ell] \Rightarrow y_{i_2} - z_{i_3} \leq \omega(\xi_1) \\ R^{N^2}(\mathbf{z}, \mathbf{x}') \Rightarrow z_{i_3} - x'_{i_4} \leq \omega(\sigma_2) \end{array} \right) \Rightarrow \left(\begin{array}{l} \psi(\ell, \mathbf{x}, \mathbf{x}')[n/\ell] \Rightarrow \\ x_{i_1} - x'_{i_4} \leq \omega(\sigma_1.\xi_1.\sigma_2) = \omega(\rho) \end{array} \right)$$

6.4 Defining the Closed Form

We finally prove that the formula $\widehat{R}(k, \mathbf{x}, \mathbf{x}')$ defined below is a closed form of R :

$$\widehat{R}(k, \mathbf{x}, \mathbf{x}') \Leftrightarrow \bigvee_{i=1}^{2N^2} (k = i \wedge R^i(\mathbf{x}, \mathbf{x}')) \vee \exists \ell \geq 1 . k = 2N^2 + \ell \wedge \psi(\ell, \mathbf{x}, \mathbf{x}') \quad (8)$$

Note that R_{fw} and R_{bw} are one-directional DB relations (see Section 4). Clearly, S_{fw} , S_{fw} , R_s , R_{fw} , and R_{bw} are PTIME-computable DB constraints, by Lemma 1 and Proposition 1. Since $\widehat{R}_{fw}(\ell, \mathbf{x}, \mathbf{x}')$ and $\widehat{R}_{bw}(\ell, \mathbf{x}, \mathbf{x}')$ are PTIME-computable formulas in the existential fragment of PA, by Theorem 1, so is the formula $\phi(\ell, \mathbf{x}, \mathbf{x}')$ in (6), and hence also $\widehat{R}(k, \mathbf{x}, \mathbf{x}')$ in (8).

Theorem 3. *Let $R(\mathbf{x}, \mathbf{x}')$ be a balanced DB constraint. Then, (8) defines a closed form of $R(\mathbf{x}, \mathbf{x}')$. Moreover, $\widehat{R}(n, \mathbf{x}, \mathbf{x}')$ is a PTIME-computable formula in the existential fragment of PA.*

By Proposition 2, the result of Theorem 3 extends to arbitrary DB relation.

Corollary 1. *Let $R(\mathbf{x}, \mathbf{x}')$ be a DB constraint. Then, its closed form is a PTIME-computable formula from the existential fragment of PA.*

7 Octagonal Relations

The class of integer octagonal constraints is defined as follows:

Definition 9. *A formula $\phi(\mathbf{x})$ is an octagonal constraint if it is a finite conjunction of terms of the form $x_i - x_j \leq a_{ij}$, $x_i + x_j \leq b_{ij}$ or $-x_i - x_j \leq c_{ij}$ where $a_{ij}, b_{ij}, c_{ij} \in \mathbb{Z}$, for all $1 \leq i, j \leq N$. A relation $R \subseteq \mathbb{Z}^{\mathbf{x}} \times \mathbb{Z}^{\mathbf{x}}$ is an octagonal relation if it can be defined by an octagonal constraint $\phi_R(\mathbf{x}, \mathbf{x}')$.*

We represent octagons as difference bounds constraints over the dual set of variables $\mathbf{y} = \{y_1, y_2, \dots, y_{2N}\}$, with the convention that y_{2i-1} stands for x_i and y_{2i} for $-x_i$, respectively. For example, the octagonal constraint $x_1 + x_2 = 3$ is represented as $y_1 - y_4 \leq 3 \wedge y_2 - y_3 \leq -3$. In order to handle the \mathbf{y} variables in the following, we define $\bar{i} = i - 1$, if i is even, and $\bar{i} = i + 1$ if i is odd. Obviously, we have $\bar{\bar{i}} = i$, for all $i \in \mathbb{N}$. We denote by $\bar{\phi}(\mathbf{y})$ the difference bounds constraint over \mathbf{y} that represents $\phi(\mathbf{x})$:

Definition 10. Given an octagonal constraint $\phi(\mathbf{x})$, $\mathbf{x} = \{x_1, \dots, x_N\}$, its difference bounds representation $\bar{\phi}(\mathbf{y})$, over $\mathbf{y} = \{y_1, \dots, y_{2N}\}$, is a conjunction of the following difference bounds constraints, where $1 \leq i, j \leq N$, $c \in \mathbb{Z}$.

$$\begin{aligned} (x_i - x_j \leq c) \in \text{atoms}(\phi) &\Leftrightarrow (y_{2i-1} - y_{2j-1} \leq c), (y_{2j} - y_{2i} \leq c) \in \text{atoms}(\bar{\phi}) \\ (-x_i + x_j \leq c) \in \text{atoms}(\phi) &\Leftrightarrow (y_{2j-1} - y_{2i-1} \leq c), (y_{2i} - y_{2j} \leq c) \in \text{atoms}(\bar{\phi}) \\ (-x_i - x_j \leq c) \in \text{atoms}(\phi) &\Leftrightarrow (y_{2i} - y_{2j-1} \leq c), (y_{2j} - y_{2i-1} \leq c) \in \text{atoms}(\bar{\phi}) \\ (x_i + x_j \leq c) \in \text{atoms}(\phi) &\Leftrightarrow (y_{2i-1} - y_{2j} \leq c), (y_{2j-1} - y_{2i} \leq c) \in \text{atoms}(\bar{\phi}) \end{aligned}$$

The following result has been proved in [3].

Lemma 5. Let $n \geq 1$ and let $R(\mathbf{x}, \mathbf{x}')$ be an octagonal relation. Then, if $R^n(\mathbf{x}, \mathbf{x}')$ is consistent, the following equivalence holds:

$$R^n(\mathbf{x}, \mathbf{x}') \Leftrightarrow \bar{R}^n(\mathbf{y}, \mathbf{y}')[\sigma], \quad \text{where} \quad \sigma = [x_i/y_{2i-1}, -x_i/y_{2i}, x'_i/y'_{2i-1}, x'_i/y'_{2i}]_{i=1}^N$$

Hence, a consistent n -th power of $R(\mathbf{x}, \mathbf{x}')$ can be computed by applying the above substitution σ on the n -th power of $\bar{R}(\mathbf{y}, \mathbf{y}')$.

Checking *-consistency We say that a relation R is **-consistent* if R^n is consistent for each $n \geq 1$. If R is not **-consistent*, we define the minimal inconsistent power of R as:

$$K_R \stackrel{\text{def}}{=} \min\{n \mid n \geq 1, R^n \text{ is inconsistent}\}$$

Lemma 6. Checking **-consistency* of R and computation of K_R can be done in PTIME.

Closed form We prove that the closed form of R can be defined as

$$\hat{R}(k, \mathbf{x}, \mathbf{x}') \Leftrightarrow \begin{cases} \hat{\bar{R}}(k, \mathbf{y}, \mathbf{y}')[\sigma] & \text{if } R \text{ is } *-consistent \\ \hat{\bar{R}}(k, \mathbf{y}, \mathbf{y}')[\sigma] \wedge k < K_R & \text{otherwise} \end{cases} \quad (9)$$

Theorem 4. Let $R(\mathbf{x}, \mathbf{x}')$ be an octagonal constraint. Then, (9) defines its closed form and moreover, it is a PTIME-computable formula in the existential fragment of PA.

8 Conclusions

We have presented a method that computes transitive closures of octagonal relations in polynomial time. This result also provides a proof of the fact that transitive closures are expressible in (the existential fragment of) Presburger arithmetic. Consequently, our result also simplifies the proof of NP-completeness of reachability checking for flat counter automata, by allowing a deterministic polynomial time reduction to the satisfiability of QFPA.

References

1. M. Bozga, C. Gîrlea, and R. Iosif. Iterating octagons. In *Proc. of TACAS*, volume 5505 of *LNCS*, pages 337–351, Berlin, Heidelberg, 2009. Springer Verlag.
2. M. Bozga, R. Iosif, and F. Konečný. Fast acceleration of ultimately periodic relations. In *Proc. of CAV*, volume 6174 of *LNCS*, pages 227–242, Berlin, Heidelberg, 2010. Springer Verlag.
3. M. Bozga, R. Iosif, and F. Konečný. Deciding conditional termination. Technical Report arXiv 1302.2762, 2013.
4. M. Bozga, R. Iosif, and F. Konečný. Safety problems are NP-complete for flat integer programs with octagonal loops. Technical Report arXiv 1307.5321, 2013.
5. M. Bozga, R. Iosif, and F. Konečný. Safety problems are NP-complete for flat integer programs with octagonal loops. In *Proc. of VMCAI*, pages 242–261, 2014.
6. M. Bozga, R. Iosif, and Y. Lakhnech. Flat parametric counter automata. *Fundamenta Informaticae*, 91(2):275–303, 2009.
7. H. Comon and Y. Jurski. Multiple counters automata, safety analysis and presburger arithmetic. In *Proc. of CAV*, volume 1427 of *LNCS*, pages 268–279, Berlin, Heidelberg, 1998. Springer Verlag.
8. H. Hojjat, R. Iosif, F. Konečný, V. Kuncak, and P. Rümmer. Accelerating interpolants. In *Proc. of ATVA*, pages 187–202, 2012.
9. H. Hojjat, F. Konečný, F. Garnier, R. Iosif, V. Kuncak, and P. Rümmer. A verification toolkit for numerical transition systems - tool paper. In *Proc. of FM*, pages 247–251, 2012.
10. A. Miné. *Weakly Relational Numerical Abstract Domains*. 2004.
11. A. Miné. The octagon abstract domain. *Higher-Order and Symbolic Computation*, 19(1):31–100, 2006.
12. C. Wang, F. Ivancic, M. K. Ganai, and A. Gupta. Deciding separation logic formulae by SAT and incremental negative cycle elimination. In *LPAR*, pages 322–336, 2005.

A Remaining Proofs from Section 3

The following technical proposition states that if R is balanced, paths in \mathcal{G}_R^n can be shifted arbitrarily in \mathcal{G}_R^n . Note that this claim doesn't hold if R is not balanced, since e.g. an extremal vertical edge $\rho = x_i^{(0)} \rightarrow x_j^{(0)}$ in \mathcal{G}_R^n has copies $\vec{\rho}^{(k)} = x_i^{(k)} \rightarrow x_j^{(k)}$ for each $k = 0, \dots, n-1$, but possibly not for $k = n$.

Proposition 8. *Let $R(\mathbf{x}, \mathbf{x}')$ be a balanced DB constraint, let $n \geq 1$, and let ρ be a path in \mathcal{G}_R^n such that $\text{positions}(\rho) = \{a, \dots, b\}$ for some $0 \leq a \leq b \leq n$. Then, $\vec{\rho}^{(k)}$ is an isomorphic path in \mathcal{G}_R^n such that $\text{positions}(\vec{\rho}^{(k)}) = \{a+k, \dots, b+k\}$, for each $k \in \{-a, \dots, n-b\}$.*

Proof of Proposition 2: It is easy to verify that $R \circ R_b^n \circ R = R^{2+n}$ for all $n \geq 0$. Then, the claim is a direct consequence of this fact. \square

B Remaining Proofs from Section 4

Proof of Lemma 2: Lemma 3 in [4] proves a general result for minimal-weight paths in weighted digraphs and the definition of a biquadratic path scheme hence refers to the

cardinality of the set of vertices instead of the number of variables N . A mapping from our notions of (minimal) paths in \mathcal{G}_R^n and path schemes to the general setting in [4] is via a technique called *zigzag automata* (weighted finite automata which can be viewed as digraphs). We refer an interested reader to [4] for the definition of zigzag automata. We only make a remark that by construction, zigzag automata of one-directional difference bounds relations have the number of control states bounded by N , from which our result follows. \square

Proof of Lemma 3: Consider the graph $\mathcal{G}_R^{N^4}$. Clearly, $\mathcal{G}_R^{N^4}$ has N^5 vertices, contains only forward edges and therefore no cycles. We will next compute $\text{min-path}(i, j, n)$, a minimal weight path from $x_i^{(0)}$ to $x_j^{(n)}$ in $\mathcal{G}_R^{N^4}$, for each $1 \leq n \leq N^4$ and $1 \leq i, j \leq N$. This computation can be done iteratively, first for $n = 0$ and $n = 1$ (note that ϵ denotes the empty path):

$$\begin{aligned} \text{min-path}(i, j, 0) &\stackrel{\text{def}}{=} \epsilon \\ \text{min-path}(i, j, 1) &\stackrel{\text{def}}{=} \begin{cases} \emptyset & \text{if } (x_i^{(0)} \xrightarrow{c_{ij}} x_j^{(1)}) \notin \mathcal{G}_R^{N^4} \\ x_i^{(0)} \xrightarrow{c_{ij}} x_j^{(1)} & \text{otherwise} \end{cases} \end{aligned}$$

and then for each $n = 2, \dots, N^4$ as:

$$\text{min-path}(i, j, n) \stackrel{\text{def}}{=} \min \left\{ \pi \cdot \pi' \mid \exists 1 \leq k \leq N \cdot \bigwedge \begin{matrix} \pi = \text{min-path}(i, k, n-1) \\ \pi' = \text{min-path}(k, j, 1) \end{matrix} \right\}$$

where the min operator is defined as

$$\min(S) \stackrel{\text{def}}{=} \begin{cases} \emptyset & \text{if } S = \emptyset \\ \text{any } \rho \in S \text{ s.t. } \forall \rho' \in S. \omega(\rho) \leq \omega(\rho') & \text{otherwise} \end{cases}$$

Correctness of this computation can be shown by induction on n . The running time of each iteration is of the order $O(N^3)$ and hence the total running time is of the order $O(N^7)$. Then, we have:

$$\Pi_{ijkpq} \neq \emptyset \Leftrightarrow \bigwedge \left\{ \begin{matrix} \text{min-path}(k, k, p) \neq \emptyset \\ \exists 0 \leq r \leq q. \text{min-path}(i, k, r) \neq \emptyset \wedge \text{min-path}(k, j, q-r) \neq \emptyset \end{matrix} \right.$$

If $\Pi_{ijkpq} \neq \emptyset$, we can define $\theta_{ijkpq} \stackrel{\text{def}}{=} \sigma \cdot \lambda^* \cdot \sigma'$ where $\lambda = \text{min-path}(k, k, p)$ and

$$\sigma \cdot \sigma' \stackrel{\text{def}}{=} \min \{ v \cdot v' \mid \exists 0 \leq r \leq q. v = \text{min-path}(i, k, r), v' = \text{min-path}(k, j, q-r) \}$$

Since $0 \leq r \leq N^4$, $\sigma \cdot \sigma'$ can be computed in $O(N^4)$ time. Since the ranges of i, j, k, p, q are $1 \leq i, j, k \leq N$, $0 \leq p \leq N$, $0 \leq q \leq N^4$, the computation of $\{\theta_{ijkpq} \mid \Pi_{ijkpq} \neq \emptyset\}$ is of the order $O(N^{12})$. \square

Proof of Theorem 1: The correctness argument follows directly from Lemma 1, Lemma 2, from the definitions of θ_{ijkpq} , S_{ij} , ϕ_{ij} , and from (4). The complexity argument follows directly from Lemma 3 and definitions of S_{ij} , ϕ_{ij} , and from (4). \square

C Remaining Proofs from Section 5

Proof of Proposition 3: (“ \Rightarrow ”) If the condition of Def. 6 holds, it follows, since every segment is subpath of ρ on positions $\{N^2, \dots, n - N^2\}$, that $\text{lb-corners}(\xi) = \emptyset$ for each $\xi \in \text{segments}(\rho)$. (“ \Leftarrow ”) To see that the converse holds, notice that every subpath θ of ρ at positions $\{N^2, \dots, n - N^2\}$ that is a long corner is also a subpath of some segment $\xi \in \text{segments}(\rho)$. Then clearly, $\text{lb-corners}(\xi) \neq \emptyset$. \square

Proof of Proposition 4: Assume that ρ is of the form $\rho = x_{i_0}^{(k_0)} \rightarrow \dots \rightarrow x_{i_m}^{(k_m)}$. Consider the following set of indices:

$$I \stackrel{\text{def}}{=} \{q \mid (x_{i_p}^{(k_p)} \rightarrow \dots \rightarrow x_{i_q}^{(k_q)}) \in \text{lb-corners}(\rho)\}$$

Since $\text{lb-corners}(\rho) \neq \emptyset$, it follows that $I \neq \emptyset$ and hence we can define $d \stackrel{\text{def}}{=} \min(I)$. By the definition of I there exists a corner $\theta \in \text{lb-corners}(\rho)$ of the form

$$\theta = x_{i_a}^{(k_a)} \rightarrow \dots \rightarrow x_{i_d}^{(k_d)}$$

for some $0 \leq a < d$. We next show that this corner is unique. By contradiction, suppose that there exists a corner $\theta' \in \text{lb-corners}(\rho)$ of the form

$$\theta' = x_{i_b}^{(k_b)} \rightarrow \dots \rightarrow x_{i_d}^{(k_d)}$$

for some $0 \leq b < d$ such that $a \neq b$. Clearly, $k_a = k_b = k_d$. Suppose that $a < b$ (the case $b < a$ is symmetric), then θ' is a subpath of θ and hence $k_a \in \{k_{a+1}, \dots, k_b = k_a, \dots, k_{d-1}\}$. Thus, θ is not basic, by the definition of a basic corner, contradiction.

Finally, we prove that $\text{extent}(\theta) = N^2 + 1$. By contradiction, suppose that e is the extent of θ such that $e > N^2 + 1$. By Proposition 11, there exists a basic corner $\theta' \in \text{b-corners}(\theta)$ such that $\theta = \eta \cdot \theta' \cdot \eta'$ for some paths η, η' and moreover, θ is of extent $N^2 + 1$ and starts at position $k' = k_0 + e - (N^2 + 1)$. Since $e > N^2 + 1$, we have $k' > k_0$. Thus, η and η' are non-empty. Consequently, θ' is of the form $x_{i_b}^{(k_b)} \rightarrow \dots \rightarrow x_{i_c}^{(k_c)}$ for some $a < b < c < d$. Clearly, $c \in I$ and since $c < d$, it follows that d is not minimal, contradiction. \square

The following proposition proves that the shape of a segment ξ determines whether the first lb-corner is right or left.

Proposition 9. *Let $n \geq 1$ and $0 \leq p \leq q \leq n$ be integers, let ρ be an extremal path in \mathcal{G}_R^n , and let $\xi \in \text{segments}(\rho, p, q)$ be a segment such that $\text{lb-corners}(\xi) \neq \emptyset$. Then, the first lb-corner of ξ is right (left) if and only if ξ starts at position p (q , respectively).*

Proof. We assume that ξ starts at position p (a proof for the starting position q is symmetric). Let θ be the first lb-corner of ξ and by contradiction, suppose that θ is left. By Proposition 4, θ has extent $N^2 + 1$. Thus, ξ is of the form

$$\xi = x_i^{(p)} \rightarrow \dots \rightarrow x_{i_1}^{(\ell+N^2+1)} \rightarrow \underbrace{\dots \rightarrow x_{i_2}^{(\ell)} \rightarrow \dots \rightarrow x_{i_3}^{(\ell+N^2+1)}}_{\theta} \rightarrow \dots \rightarrow x_j^{(k)}$$

for some $1 \leq i_1, i_2, i_3 \leq N$ and some integers k, ℓ . Without loss of generality, we assume that $x_{i_2}^{(\ell)}$ is the first vertex of θ at position ℓ . Since $\text{positions}(\xi) \subseteq \{p, p+1, \dots, q\}$, it follows that $\ell \geq p$. Consequently, the subpath $x_i^{(p)} \rightarrow \dots \rightarrow x_{i_1}^{(\ell+N^2+1)}$ must traverse the position ℓ . Let $x_{i_0}^{(\ell)}$ be the last vertex of this subpath that is at position ℓ . Then, ξ can be written as

$$\xi = x_i^{(p)} \rightarrow \dots \rightarrow x_{i_0}^{(\ell)} \rightarrow \underbrace{\dots \rightarrow x_{i_1}^{(\ell+N^2+1)} \rightarrow \dots \rightarrow x_{i_2}^{(\ell)}}_{\theta'} \rightarrow \dots \rightarrow x_{i_3}^{(\ell+N^2+1)} \rightarrow \dots \rightarrow x_j^{(k)}$$

Note that θ' is a lb-corner which contradicts the fact that θ is the first lb-corner. \square

Proposition 10. (Shape of basic corners) *Let ρ be a basic right (left) corner such that $\text{extent}(\rho) \geq 1$. Then, ρ is of the form $\rho = \pi.\rho'. where π is a forward (backward) edge, π' is a backward (forward) edge, and ρ' is a right (left) corner such that $\text{extent}(\rho') = \text{extent}(\rho) - 1$.$*

Proof. We assume that ρ is a basic right corner (a proof for a basic right corner is symmetric). Since $\text{extent}(\rho) \geq 1$, ρ consists of at least two edges and hence is of the form $\rho = \pi.\rho'. for some edges π, π' and a path ρ' . Let k be the position where ρ starts and ends. Since ρ is a right corner, then $\text{positions}(\rho) = \{k, \dots, k + \text{extent}(\rho)\}$. Since ρ is basic, it follows that $k \notin \text{positions}(\rho')$. Hence, π (π') ends (starts) at position $k+1$ and $\text{positions}(\rho') = \{k+1, \dots, k + \text{extent}(\rho)\}$. Consequently, π (π') is forward (backward), ρ' is a right corner that starts at position $k+1$, and $\text{extent}(\rho') = (k + \text{extent}(\rho)) - (k+1) = \text{extent}(\rho) - 1$. $\square$$

Proposition 11. (Basic subcorners) *Let ρ be a right corner of extent e that starts at position k and let $d \in \{1, \dots, e\}$ be an integer. Then, there exists a right corner $\rho' \in \text{b-corners}(\rho)$ of extent d such that ρ' starts at position $k+e-d$.*

Proof. Clearly, $\text{positions}(\rho) = \{k, \dots, k+e\}$. Let u and w be the first and last vertex of ρ , respectively, and let v be any vertex of ρ at position $k+e$. Then, ρ is of the form $u \rightarrow \dots \rightarrow v \rightarrow \dots \rightarrow w$. Consider the subpath $u \rightarrow \dots \rightarrow v$ and let u' be the its last vertex at position $k+e-d$. Consider the subpath $v \rightarrow \dots \rightarrow w$ and let w' be its first vertex at position $k+e-d$. We define $\rho' \stackrel{\text{def}}{=} u' \rightarrow \dots \rightarrow v \rightarrow \dots \rightarrow w'$. It is easy to see that ρ' is a basic corner whose positions range between $k+e-d$ and $k+e$ and hence the extent of ρ' is d . \square

Proposition 12. *For each path ρ , $\text{lb-corners}(\rho) = \emptyset$ if and only if $\text{l-corners}(\rho) = \emptyset$.*

Proof. (“ \Leftarrow ”) Obvious, by the definition of $\text{lb-corners}(\rho)$. (“ \Rightarrow ”) We prove the converse. Suppose there exists $\theta \in \text{l-corners}(\rho)$. Clearly, $\text{extent}(\theta) > N^2$. By Proposition 11, there exists $\theta' \in \text{b-corners}(\theta) \subseteq \text{b-corners}(\rho)$ such that $\text{extent}(\theta') = \text{extent}(\theta) > N^2$. Hence, $\theta' \in \text{b-corners}(\rho)$ and $\theta' \in \text{l-corners}(\rho)$. Consequently, $\theta' \in \text{lb-corners}(\rho)$, by the definition of $\text{lb-corners}(\rho)$. \square

Proposition 13. *Let μ be a repeating path such that $\text{card}(\text{positions}(\mu)) \leq N^2 + 1$. Then, $\text{lb-corners}(\mu^k) = \emptyset$ for all $k \geq 1$.*

Proof. We assume that μ is a forward path (the proof for μ being backward is symmetric). By contradiction, suppose that for some $k \geq 1$, μ^k contains a long right corner θ (the argument for a left corner is symmetric). Then, θ is of the form

$$x_{j_0}^{(\ell)} \rightarrow \dots \rightarrow x_{j_1}^{(\ell+N^2+1)} \rightarrow \dots \rightarrow x_{j_2}^{(\ell)}$$

for some $1 \leq j_0, j_1, j_2 \leq N$ and for some integer ℓ . Clearly, there exist integers $1 \leq p \leq q \leq r \leq k$ such that $x_{j_0}^{(\ell)}, x_{j_1}^{(\ell+N^2+1)}, x_{j_2}^{(\ell)}$, respectively) is traversed in the p -th (q -th, r -th, respectively) copy of μ in μ^k . Since μ^k consists of k copies of a forward repeating path μ , it follows that the r -th copy traverses also $x_{j_1}^{(\ell+N^2+1+(r-q) \cdot \|\mu\|)}$, in addition to $x_{j_2}^{(\ell)}$. Since $(r-q) \cdot \|\mu\| \geq 0$, it follows that $\{\ell, \dots, \ell + N^2 + 1\} \subseteq \text{positions}(\mu)$, which contradicts the assumption that $\text{card}(\text{positions}(\mu)) \leq N^2 + 1$. Hence, the claim holds. \square

Proposition 14. *Let ρ be a right corner of the form $\rho = x_i^{(k)} \rightarrow \dots \rightarrow x_j^{(k)}$ and let $1 \leq d \leq \text{extent}(\rho)$. Then, ρ has a subpath ρ' of the form $\rho' = x_i^{(k)} \rightarrow \dots \rightarrow x_p^{(k+d)}$ for some $1 \leq p \leq N$ such that $\text{positions}(\rho') = \{k, \dots, k+d\}$.*

Proof. We choose $x_p^{(k+d)}$ as the first vertex of ρ at position $k+d$ and define ρ' as the subpath of ρ that starts at $x_i^{(k)}$ and ends at $x_p^{(k+d)}$. Since ρ is a corner and ρ' its subpath, then $\text{positions}(\rho') \subseteq \text{positions}(\rho) = \{k, \dots, k + \text{extent}(\rho)\}$. By the choice of $x_p^{(k+d)}$ and by the definition of ρ' , we have that $k+d \in \text{positions}(\rho')$ and $k+d+1 \notin \text{positions}(\rho')$. Consequently, $\text{positions}(\rho') = \{k, \dots, k+d\}$. \square

Proposition 15. (Corner shortening / decomposition) *Let ρ be a right (left) corner such that $\text{extent}(\rho) = N^2$. Then, there exists a compatible right (left) corner ρ' such that either (i) $\text{extent}(\rho') < N^2$ or (ii) $\text{extent}(\rho') = N^2$ and ρ' has subpaths $\eta, \mu, \tau, \mu', \eta'$ such that $\rho' = \eta \cdot \mu \cdot \tau \cdot \mu' \cdot \eta'$, μ is a forward (backward) repeating path, μ' is a backward (forward) repeating path, τ is a right (left) corner, $\|\eta\| = \|\eta'\|$, $1 \leq \|\mu\| = \|\mu'\| \leq N^2$, and $\omega(\mu) + \omega(\mu') < 0$. Moreover, for all $k \geq 0$, $\eta \cdot \mu^k \cdot \tau \cdot \mu'^k \cdot \eta'$ is a right (left) corner and $\text{lb-corners}(\eta \cdot \mu^k) = \text{lb-corners}(\mu'^k \cdot \eta') = \emptyset$.*

Proof. Suppose that ρ is a right corner (a proof for left corner is symmetric). First, we assume that ρ starts at position 0. Then, since $\text{extent}(\rho) = N^2$ there exist paths ρ_1, ρ_2 such that $\rho = \rho_1 \cdot \rho_2$, and ρ_1 (ρ_2) starts at position 0 (N^2) and ends at position N^2 (0). Next, ρ can be decomposed into the following form:

$$\rho = (\underbrace{\xi_1 \cdot \xi_2 \dots \xi_{N^2}}_{\rho_1}) \cdot (\underbrace{\zeta_{N^2} \dots \zeta_2 \zeta_1}_{\rho_2})$$

where for each $1 \leq m \leq N^2$, ξ_m and ζ_m are of the form:

$$\xi_m = x_{i_{m-1}}^{(m-1)} \rightarrow \dots \rightarrow x_{i_m}^{(m)} \quad \zeta_m = x_{j_m}^{(m)} \rightarrow \dots \rightarrow x_{j_{m-1}}^{(m-1)}$$

(note that $x_{i_{N^2}}^{(N^2)} = x_{j_{N^2}}^{(N^2)}$) and moreover:

$$\text{positions}(\xi_m), \text{positions}(\zeta_m) \subseteq \{m-1, \dots, N^2\} \quad (10)$$

The property (10) can be achieved as follows. First, $x_{i_{N^2}}^{(N^2)}$ can be picked as the first vertex of ρ at position N^2 . This induces the above partition of ρ into $\rho = \rho_1 \cdot \rho_2$. Then, for each $1 \leq m < N^2$, the vertex $x_{i_m}^{(m)}$ is chosen as the immediate successor of the last vertex of ρ_1 at position $m-1$. Symmetrically, for each $1 \leq m < N^2$, the vertex $x_{j_m}^{(m)}$ is chosen as the immediate predecessor of the first vertex of ρ_2 at position $m-1$.

Next, we define the sequence $\{(i_m, j_m)\}_{m=0}^{N^2}$. By the pigeonhole principle, there exist integers $0 \leq a < b \leq N^2$ such that $(i_a, j_a) = (i_b, j_b)$. We define:

$$\mu \stackrel{\text{def}}{=} \xi_{a+1} \xi_{a+2} \dots \xi_b \quad \mu' \stackrel{\text{def}}{=} \zeta_b \zeta_{b-1} \dots \zeta_{a+1}$$

Note that μ (μ') is a forward (backward) repeating path of the form

$$\mu = x_{i_a}^{(a)} \rightarrow \dots \rightarrow x_{i_a}^{(b)} \quad \mu' = x_{j_a}^{(b)} \rightarrow \dots \rightarrow x_{j_a}^{(a)}$$

and $\|\mu\| = \|\mu'\| = b-a \geq 1$. Clearly, there exist (possibly empty) paths η, τ, η' such that $\rho = \eta \cdot \mu \cdot \tau \cdot \mu' \cdot \eta'$. Note that η, τ, η' are of the form:

$$\eta = \xi_1 \xi_2 \dots \xi_a \quad \tau = \xi_{b+1} \dots \xi_{N^2} \cdot \zeta_{N^2} \dots \zeta_{b+1} \quad \eta' = \zeta_a \zeta_{a-1} \dots \zeta_1$$

Note that τ traverses the vertex $x_{i_{N^2}}^{(N^2)}$. Clearly, we have $\|\eta\| = \|\eta'\|$ and $1 \leq \|\mu\| = b-a = \|\mu'\| \leq N^2$. It is easy to verify that (10) entails that $\tau, \mu^k \cdot \tau \cdot \mu'^k$, and $\eta \cdot \mu^k \cdot \tau \cdot \mu'^k \cdot \eta'$ are right corners, for all $k \geq 0$.

If $N^2 \in \text{positions}(\eta \cdot \tau \cdot \eta')$ and $\omega(\mu) + \omega(\mu') \geq 0$, we assign $\rho \leftarrow \eta \cdot \tau \cdot \eta'$ and apply the above construction again. We iterate this construction until either $N^2 \notin \text{positions}(\eta \cdot \tau \cdot \eta')$ or $\omega(\mu) + \omega(\mu') < 0$. Note that this iteration terminates, since at each step, $\eta \cdot \tau \cdot \eta'$ traverses the position N^2 strictly less times than ρ does; this is because $\|\mu\| = \|\mu'\| = b-a \geq 1$ and hence, all vertices in τ (including $x_{i_{N^2}}^{(N^2)}$) are shifted by $b-a$ when constructing $\eta \cdot \tau \cdot \eta'$, hence the vertex $x_{i_{N^2}}^{(N^2)}$ in τ will become $x_{i_{N^2}}^{(N^2-(b-a))}$ in $\eta \cdot \tau \cdot \eta'$. Then, we define ρ' to be the path from the last iteration. Either we have (i) $\text{extent}(\rho') < N^2$ (this happens when $N^2 \notin \text{positions}(\rho')$) or (ii) $\omega(\mu) + \omega(\mu') < 0$. We next prove that in case (ii), $\text{lb-corners}(\eta \cdot \mu^k) = \text{lb-corners}(\mu'^k \cdot \eta') = \emptyset$, for all $k \geq 0$.

It follows from (10) and from the definition of μ that

$$\text{positions}(\eta \cdot \mu^k \cdot \tau \cdot \mu'^k \cdot \eta') \subseteq \{0, \dots, N^2 + \|\mu\| \cdot k\} \quad (11)$$

for all $k \geq 0$ and hence $\eta \cdot \mu^k \cdot \tau \cdot \mu'^k \cdot \eta'$ is a right corner for all $k \geq 0$. Next, we show that $\text{lb-corners}(\eta \cdot \mu^k) = \emptyset$ for all $k \geq 0$ (a proof of $\text{lb-corners}(\mu'^k \cdot \eta') = \emptyset$ is symmetric). Since $\text{extent}(\rho) = N^2$ and $\eta, \eta \cdot \mu$ are subpaths of ρ , it follows that $\text{lb-corners}(\rho) = \text{lb-corners}(\eta \cdot \mu) = \text{lb-corners}(\eta) = \emptyset$. By (10) and by the definition of μ , we have $\text{positions}(\mu) \subseteq$

$\{0, \dots, N^2\}$ and hence $\text{card}(\text{positions}(\mu)) \leq N^2 + 1$. Thus, by Proposition 13, $\text{lb-corners}(\mu^k) = \emptyset$. By contradiction, suppose that $\text{lb-corners}(\eta, \mu^k) \neq \emptyset$ for some $k \geq 0$ and let θ be a corner from $\text{lb-corners}(\eta, \mu^k)$. Since $\text{lb-corners}(\eta) = \text{lb-corners}(\mu^k) = \emptyset$, it follows that θ must traverse edges of both η and μ^k . Hence, $\theta \in \text{positions}(\theta)$. Since $\text{lb-corners}(\eta, \mu) = \emptyset$, it follows that $k \geq 2$. By the definition of μ , we have that μ^k starts at position a and

$$\text{positions}(\mu^k) \subseteq \{a, \dots, N^2 + (b-a) \cdot \|\mu\|\} \quad (12)$$

Consequently, θ ends at some position $c \geq a$. We prove that θ is a left corner. By contradiction, suppose that θ is a right corner. By the definition of right corners, it cannot be that $c > a$. By (12), we have $c \geq a$. Hence we infer that $c = a$. Thus, θ ends (and consequently also starts) at position a . Since θ contains at least one edge of both η and μ^k , the position a is traversed at least three times, which contradicts the definition of a basic corner. Hence, θ is a left corner. Since $\text{positions}(\eta) \subseteq \{0, \dots, N^2\}$ and since θ contains at least one edge of η , it follows that θ starts at some position $c \in \{0, \dots, N^2\}$. Then, we infer, from the fact that θ is left and from (11), that $\text{positions}(\theta) \subseteq \{0, \dots, c\}$. Since $c \leq N^2$, it follows that θ is short. Hence, $\text{lb-corners}(\eta, \mu^k) = \text{lb-corners}(\mu^k, \eta') = \emptyset$, for all $k \geq 0$.

Finally, note that the arguments above (for 0 as the initial position of ρ) don't actually depend on the initial position of ρ and hence generalize to corners that start at arbitrary position. \square

Proof of Lemma 4: The result is obtained by chaining Proposition 10 and Lemma 4. \square

D Remaining Proofs from Section 6

Proof of Proposition 5: Let $n \geq 1$ and suppose that θ is a short right corner in \mathcal{G}_R^n , i.e. of the form $\theta = x_i^{(k)} \rightarrow \dots \rightarrow x_j^{(k)}$ such that $\text{positions}(\theta) \subseteq \{k, \dots, k + N^2\}$. By Proposition 8, there is an isomorphic corner $\theta' = x_i^{(0)} \rightarrow \dots \rightarrow x_j^{(0)}$ such that $\text{positions}(\theta') \subseteq \{0, \dots, N^2\}$, i.e. θ' appears in \mathcal{G}_R^n . Hence, by Lemma 1, we have:

$$\exists \mathbf{x}' . R^{N^2}(\mathbf{x}, \mathbf{x}') \Rightarrow x_i - x_j \leq \omega(\theta)$$

we infer (due to the conjunct S_{fw} of R_s) that $R_s \Rightarrow x_i - x_j \leq \omega(\theta)$ and consequently, that \mathcal{G}_{R_s} contains an edge $x_i \xrightarrow{c} x_j$ for some $c \leq \omega(\theta)$.

Next, we prove that $\theta' \stackrel{\text{def}}{=} x_i^{(k)} \xrightarrow{c} x_j^{(k)}$ is an edge in $\mathcal{G}_{R_s}^n$. If $k = n$, θ contains only vertical edges and consequently, \mathcal{G}_R has an implied edge $x'_i \xrightarrow{c} x'_j$, $c \leq \omega(\theta)$, and hence \mathcal{G}_R^n (and consequently $\mathcal{G}_{R_s}^n$) has an edge $x_i^{(n)} \xrightarrow{c} x_j^{(n)}$. If $k < n$, since $\mathcal{G}_{R_s}^n$ consists of n identical copies of \mathcal{G}_{R_s} and $\theta' = x_i \xrightarrow{c} x_j$ is an edge in S_{fw} , it follows that $\mathcal{G}_{R_s}^n$ has an edge $x_i^{(\ell)} \xrightarrow{c} x_j^{(\ell)}$ for each $\ell = 0, \dots, k-1$. Finally, observe that the conjunct S_{bw} of R_s allows us to make a symmetric argument about short left corners. \square

Proof of Proposition 6: Let $n \geq 1$ and ξ be an extremal path in \mathcal{G}_R^n , i.e. a path of the form $x_i^{(p)} \rightarrow \dots \rightarrow x_j^{(q)}$ for some $p, q \in \{0, n\}$. We next analyze the four cases induced by values of p, q . If $p = q = 0$, then ξ is a (short) corner and hence, by Proposition 5:

$$S_{fw} \Rightarrow x_i - x_j \leq \omega(\xi) \quad (13)$$

If $p = q = n$, we symmetrically obtain that $S_{bw} \Rightarrow x_i - x_j \leq \omega(\xi)$.

Next, suppose that $p = 0$ and $q = n$, i.e. that ξ is forward. It is not difficult to show that ξ can be written as a sequence of (i) right corners and (ii) forward edges. For each right corner, there is a compatible vertical edge in $\mathcal{G}_{R_s}^n$, by Proposition 5. Thus, we obtain path ξ' in $\mathcal{G}_{R_s}^n$ that is compatible with ξ and is a sequence of (i) vertical edges and (ii) forward edges. Next, it is easy to see that every subpath of ξ' which follows the pattern *(vertical-edge)*.forward-edge* can be replaced by a single (implied) forward edge. Similarly, the suffix of ξ' of the form *forward-edge.(vertical-edge)** can be replaced by an implied forward edge. Thus, there exists a path ξ'' in $\mathcal{G}_{R_s}^n$ which is compatible with ξ' (and hence with ξ) and which is a sequence of forward edges only. Then, it is not hard to show that ξ'' appears also in $\mathcal{G}_{R_{fw}}^n$. Since ξ'' as an extremal path in $\mathcal{G}_{R_{fw}}^n$, we infer, by Lemma 1, that

$$R_{fw}^n(\mathbf{x}, \mathbf{x}') \Rightarrow x_i - x'_j \leq \omega(\xi'') \leq \omega(\xi)$$

Supposing that $\widehat{R}_{fw}(\ell, \mathbf{x}, \mathbf{x}')$ is a closed forms of R_{fw} , we obtain that

$$\widehat{R}_{fw}(\ell, \mathbf{x}, \mathbf{x}') [n/\ell] \Rightarrow x_i - x_j \leq \omega(\xi) \quad (14)$$

If $p = n$ and $q = 0$, i.e. ξ is backward, we can make a symmetric argument. Then, it follows from the above reasoning (from Eq. (13), (14), and the two symmetric cases) that $\phi(\ell, \mathbf{x}, \mathbf{x}')$ encodes all extremal paths in \mathcal{G}_R^n that have no long corners. \square

Proof of Proposition 7: Let $n \geq 1$ and ρ be a extremal path in $\mathcal{G}_R^{2N^2+n}$. By Theorem 2, we can assume, without loss of generality, that ρ is normalized. Recall from Definition 8 that ρ can be written as

$$\rho = \sigma_1 \cdot \xi_1 \dots \sigma_m \cdot \xi_m \cdot \sigma_{m+1} \quad (15)$$

for some $m \geq 1$ where $\{\xi_1, \dots, \xi_m\} = \text{segments}(\rho, N^2, N^2 + n)$ and every segment ξ_k , $1 \leq k \leq m$, is of the form:

$$\xi_k = x_{i_k}^{(N^2+p_k)} \rightarrow \dots \rightarrow x_{j_k}^{(N^2+q_k)} \text{ where } 1 \leq i_k, j_k \leq N, p_k, q_k \in \{0, n\}$$

Since $\text{positions}(\xi_k) \subseteq \{N^2, \dots, N^2 + n\}$, there exists, by Proposition 8, an isomorphic path ξ'_k in $\mathcal{G}_R^{2N^2+n}$ of the form $\xi'_k = x_{i_k}^{(p_k)} \rightarrow \dots \rightarrow x_{j_k}^{(q_k)}$ such that $\text{positions}(\xi'_k) \subseteq \{0, \dots, n\}$, i.e. ξ'_k appears in \mathcal{G}_R^n . Since, $p_k, q_k \in \{0, n\}$, it follows that ξ'_k is extremal in \mathcal{G}_R^n . Since ξ' contain no long corners, by definition 6, and is an extremal path in \mathcal{G}_R^n , we infer, by Proposition 6, that ξ' (and thus ξ) is encoded in $\phi(\ell, \mathbf{y}, \mathbf{z})$.

It follows from Definition 8 that each σ_k , $1 \leq k \leq m+1$, appears in the first N^2 or in the last N^2 copies of \mathcal{G}_R , formally:

$$\text{positions}(\sigma_k) \subseteq \{0, \dots, N^2\} \text{ or } \text{positions}(\sigma_k) \subseteq \{N^2 + n, \dots, 2N^2 + n\}$$

It is easy to see that σ_k , $1 \leq i \leq m+1$, is encoded either in $R^{N^2}(\mathbf{x}, \mathbf{y})$ or in $R^{N^2}(\mathbf{z}, \mathbf{x}')$.

E.g., suppose that $\rho = \sigma_1.\xi_1.\sigma_2.\xi_2.\sigma_3$ where

$$\begin{aligned} \sigma_1 &= x_{i_1}^{(0)} \rightarrow \dots \rightarrow x_{i_2}^{(N^2)} & \xi_1 &= x_{i_2}^{(N^2)} \rightarrow \dots \rightarrow x_{i_3}^{(N^2+n)} \\ \sigma_2 &= x_{i_3}^{(N^2+n)} \rightarrow \dots \rightarrow x_{i_4}^{(N^2+n)} & \xi_2 &= x_{i_4}^{(N^2+n)} \rightarrow \dots \rightarrow x_{i_5}^{(N^2)} \\ \sigma_3 &= x_{i_5}^{(N^2)} \rightarrow \dots \rightarrow x_{i_6}^{(0)} \end{aligned}$$

Then, $\sigma_1, \xi_1, \sigma_2, \xi_2, \sigma_3$ are encoded as:

$$\begin{aligned} R^{N^2}(\mathbf{x}, \mathbf{y}) &\Rightarrow x_{i_1} - y_{i_2} \leq \omega(\sigma_1) \wedge \phi(\ell, \mathbf{y}, \mathbf{z})[n/\ell] \Rightarrow y_{i_2} - z_{i_3} \leq \omega(\xi_1) \wedge \\ R^{N^2}(\mathbf{z}, \mathbf{x}') &\Rightarrow z_{i_3} - z_{i_4} \leq \omega(\sigma_2) \wedge \phi(\ell, \mathbf{y}, \mathbf{z})[n/\ell] \Rightarrow z_{i_4} - y_{i_5} \leq \omega(\xi_2) \wedge \\ R^{N^2}(\mathbf{x}, \mathbf{y}) &\Rightarrow y_{i_5} - x_{i_6} \leq \omega(\sigma_3) \end{aligned}$$

By transitivity, we infer that ρ is encoded in Ψ :

$$\Psi(\ell, \mathbf{x}, \mathbf{x}') [n/\ell] \Rightarrow x_{i_1} - x_{i_6} \leq \omega(\sigma_1) + \omega(\xi_1) + \omega(\sigma_2) + \omega(\xi_2) + \omega(\sigma_3) = \omega(\rho)$$

In general, for a path ρ of the form (15), we obtain $2m+1$ constraints implied by $R^{N^2}(\mathbf{x}, \mathbf{y}) \wedge \phi(\ell, \mathbf{y}, \mathbf{z}) \wedge R^{N^2}(\mathbf{z}, \mathbf{x}')$, one constraint for each path $\tau_1, \xi_1, \dots, \tau_m, \xi_m, \tau_{m+1}$. Then, by transitivity of these constraints, we infer that ρ is encoded in Ψ . \square

Proof of Theorem 3: We need to prove that for all $n \geq 1$:

$$\widehat{R}(k, \mathbf{x}, \mathbf{x}') [n/k] \Leftrightarrow R^n(\mathbf{x}, \mathbf{x}') \quad (16)$$

It is easy to verify that (16) holds for each $n = 1, \dots, 2N^2$. Similarly, one can check that for every $n > 2N^2$, (16) is equivalent to

$$\Psi(\ell, \mathbf{x}, \mathbf{x}') [n - 2N^2/\ell] \Leftrightarrow R^n(\mathbf{x}, \mathbf{x}')$$

Equivalently, we can prove that for all $n \geq 1$:

$$\Psi(\ell, \mathbf{x}, \mathbf{x}') [n/\ell] \Leftrightarrow R^{2N^2+n}(\mathbf{x}, \mathbf{x}') \quad (17)$$

By Proposition 7, every extremal path in $\mathcal{G}_R^{2N^2+n}$ is encoded in $\Psi(\ell, \mathbf{x}, \mathbf{x}') [n/k]$. Consequently, by Lemma 1, we have that $\Psi(\ell, \mathbf{x}, \mathbf{x}') [n/k] \Rightarrow R^{2N^2+n}(\mathbf{x}, \mathbf{x}')$. It remains to show that the “ \Leftarrow ” implication of (17) holds too. The intuition behind this is that for the strengthened relation R_s , one can prove that for all $n \geq 0$:

$$R^{2N^2+n} \Leftrightarrow \underbrace{R^{N^2}}_{\text{prefix}} \circ R_s^n \circ \underbrace{R^{N^2}}_{\text{suffix}} \quad (18)$$

Intuitively, this is because the strengthening of R with $\exists \mathbf{x}'. R^{N^2}$ (with $\exists \mathbf{x}. R^{N^2}$) gives only constraints which are already implied by the above suffix (prefix) R^{N^2} . By the definition of R_s , we have that $R_s \Rightarrow S_{fw}(\mathbf{x})$ and $R_s \Rightarrow S_{bw}(\mathbf{x}')$. It follows from the definition of R_{fw} that $R_s^n \Rightarrow R_{fw}^n$. Similarly, $R_s^n \Rightarrow R_{bw}^n$. Thus, we have:

$$\begin{aligned} R^{N^2} \circ R_s^n \circ R^{N^2} &\Rightarrow R^{N^2} \circ [R_{fw}^n \wedge R_{bw}^n \wedge S_{fw}(\mathbf{x}) \wedge S_{bw}(\mathbf{x}')] \circ R^{N^2} \\ &\Leftrightarrow \Psi(\mathbf{x}, \mathbf{x}', \ell)[n/\ell] \end{aligned} \quad (19)$$

The last equivalence above is by the definition of Ψ . Combining (18) with (19), we infer that the “ \Leftarrow ” implication of (17) also holds. \square

E Remaining Proofs from Section 7

Proof of Lemma 5: See Eq. (4.16) in the proof of Proposition 15 in [3]. \square

Proof of Lemma 6: If a relation R is not $*$ -consistent, there exists a computable exponential upper bound B on K_R of the order $2^{O(\|R\|_2)}$ (see Lemma 20 in [4]). Since the n -th power of an octagonal relation can be computed in $O(\|R\|_2 \cdot \log_2 n)$ time (see Lemma 18 in [4]), one can find K_R by performing a binary search on the interval $\{1, \dots, B\}$ and use the fact that the powers of R can be computed by fast exponentiation. \square

Proof of Theorem 4: First, suppose that R is $*$ -consistent. Then, for each $n \geq 1$, we have that $R^n(\mathbf{x}, \mathbf{x}')$ is consistent and hence:

$$\begin{aligned} \widehat{R}(k, \mathbf{x}, \mathbf{x}')[n/k] &\Leftrightarrow \widehat{\bar{R}}(k, \mathbf{y}, \mathbf{y}')[n/k][\sigma] && \text{(by (9))} \\ &\Leftrightarrow \bar{R}^n(\mathbf{y}, \mathbf{y}')[\sigma] \\ &\Leftrightarrow R^n(\mathbf{x}, \mathbf{x}') && \text{(by Lemma 5)} \end{aligned}$$

Next, suppose that R is not $*$ -consistent. For each $1 \leq n < K_R$, we have that $R^n(\mathbf{x}, \mathbf{x}')$ is consistent and hence:

$$\begin{aligned} \widehat{R}(k, \mathbf{x}, \mathbf{x}')[n/k] &\Leftrightarrow \widehat{\bar{R}}(k, \mathbf{y}, \mathbf{y}')[n/k][\sigma] \wedge n < K_R && \text{(by (9))} \\ &\Leftrightarrow \bar{R}^n(\mathbf{y}, \mathbf{y}')[\sigma] \wedge \mathbf{true} \\ &\Leftrightarrow R^n(\mathbf{x}, \mathbf{x}') && \text{(by Lemma 5)} \end{aligned}$$

For each $n \geq K_R$, we obtain:

$$\begin{aligned} \widehat{R}(k, \mathbf{x}, \mathbf{x}')[n/k] &\Leftrightarrow \widehat{\bar{R}}(k, \mathbf{y}, \mathbf{y}')[n/k][\sigma] \wedge n < K_R && \text{(by (9))} \\ &\Leftrightarrow \mathbf{false} \end{aligned}$$

Hence, in all cases, (9) defines a closed form of $R(\mathbf{x}, \mathbf{x}')$. The fact that $\widehat{R}(k, \mathbf{x}, \mathbf{x}')$ is PTIME-computable in the existential fragment of Presburger arithmetic follows immediately from Lemma 6 and Corollary 1. \square